

FAQs: Microsoft Virtual Academy

Windows IT Pro

A PENTON PUBLICATION

NOVEMBER 2012 | WINDOWSITPRO.COM | WE'RE IN IT WITH YOU

Secure the Private Cloud

Plus >>

**Best Practices for
Group Policy Design**

**Business Benefits of
Unified Communications**

**10 Steps to Migrate to
Configuration Manager 2012**

**Mark Minasi's
Windows
Power Tools**



**Automate
PowerShell Reports**

1&1 DYNAMIC CLOUD SERVER

Our data centers offer top security, Cisco firewall protection and maximum uptime. With more than 20 years experience and an extensive server range, we know what IT professionals need. Get full root access for complete control. We are a strong global company with 3 billion dollars in annual revenue and over 6,000 employees worldwide.



LIFETIME DISCOUNT
50% OFF
INCLUDING CONFIGURATIONS,
NO SETUP FEE

1&1 DYNAMIC CLOUD SERVER

A fully flexible server for a range of requirements including applications, databases, gaming and much more!

- Independently configure CPU, RAM, and storage
- Accurate and fair: Control costs with pay-per-configuration and hourly billing
- Up to 6 Cores, 24 GB RAM, 800 GB storage
- 2000 GB of traffic included free
- Parallels® Plesk Panel 11 for unlimited domains, reseller ready
- Up to 99 virtual machines with different configurations under one contract
- No setup fee
- 24/7 phone and e-mail support

\$24.99 per month* ~~\$49.99~~ per month



MAXIMUM FLEXIBILITY

Independently adjust CPU cores, RAM and hard disk space and add up to 99 virtual machines. We offer cost transparency through hourly billing.



MAXIMUM SECURITY

Redundant storage and mirrored processing units reliably protect your server against any failure



PARALLELS PLESK® PANEL 11

for unlimited domains



FULL ROOT ACCESS

The control and functionality of a root server with dedicated resources



INCLUDED TRAFFIC

2000 GB included

Parallels®
Plesk Panel

SUSE



www.1and1.com



*Offer valid for a limited time only. Lifetime 50% off applies to base fee and configurations. Base configuration includes 1 processor core, 1 GB RAM, 100 GB storage. This offer applies to new contracts only, 12 month minimum contract term. Other terms and conditions may apply. Visit www.1and1.com for full promotional offer details. Program and pricing specifications and availability subject to change without notice. 1&1 and the 1&1 logo are trademarks of 1&1 Internet, all other trademarks are the property of their respective owners. © 2012 1&1 Internet. All rights reserved.



CLOUD TRANSFORMS IT

EMC²

COVER STORY ▼

Secure the Private Cloud 38

— John Howie

Private clouds bring many benefits to an organization. Don't let a lack of security practices get in the way of cloud adoption.

Features

56 Best Practices for Group Policy Design

Darren Mar-Elia

71 Business Benefits of Unified Communications

Nathan Winters

84 10 Steps to Migration: Configuration Manager 2012

Peter Daalmans

Special Features

36 Microsoft Releases Windows Server 2012

96 Microsoft Windows 8 Arrives

Interact

27 Reader to Reader

29 Ask the Experts

In Every Issue

9 IT Community Forum

116 Ctrl+Alt+Del

117 Advertiser Directory

117 Directory of Services

117 Vendor Directory

Chat with Us



Facebook



Twitter



LinkedIn

Columns

6 IT Pro Perspectives

IT Pros Are from Mars, Marketing Pros Are from Venus

Sean Deuby



11 Need to Know

Hands-On Time with Window Phone 8 and Why Windows RT Is Important

Paul Thurrott



17 Windows Power Tools

Automating PowerShell Reports, Part 2

Mark Minasi



20 Top 10

FAQs About Microsoft Virtual Academy

Michael Otey



23 Enterprise Identity

Building Your Identity Bridge to the Cloud

Sean Deuby



Products

98 New & Improved

102 Hard Disk Manager 12 Professional

Eric B. Rux

107 Ericom AccessNow 2.0

Russell Smith

111 Industry Bytes

Editorial

Editorial Director:
Megan Keller
Editor-in-Chief:
Amy Eisenberg
Senior Technical Director:
Michael Otey
Technical Director:
Sean Deuby
Senior Technical Analyst:
Paul Thurrott
Custom Group Editorial Director:
Dave Bernard
Exchange & Outlook:
Brian Winstead
Systems Management,
Networking, Hardware:
Jason Bovberg
Scripting:
Blair Greenwood
Security, Virtualization:
Amy Eisenberg
SharePoint, Active Directory:
Caroline Marwitz
SQL Server, Developer Content:
Megan Keller
Managing Editor:
Lavon Peters
Assistant Managing Editor:
Rachel Koon
Editorial SEO Specialist:
Jayleen Heft

Senior Contributing Editors

David Chernicoff, Mark Minasi,
Tony Redmond, Paul Robichaux,
Mark Russinovich, John Savill

Contributing Editors

Alex K. Angelopoulos, Michael Dragone,
Jeff Felling, Brett Hill, Dan Holme,
Darren Mar-Elia, Eric B. Rux,
William Sheldon, Curt Spanburgh,
Bill Stewart, Orin Thomas,
Douglas Toombs, Ethan Wilansky

Art & Production

Production Director: Linda Kirchgesler
Senior Graphic Designer: Matt Wiebe

Advertising Sales

Publisher: Peg Miller
Key Account Director:
Chrissy Ferraro • 970-203-2883
Account Executives:
Barbara Ritter • 858-367-8058
Cass Schulz • 858-357-7649

Client Services

Sales Operation Manager:
Patti McKinzie • 970-613-4922
Senior Client Services Manager:
Michelle Andrews • 970-613-4964
Client Services Manager:
Glenda Vaught • 970-203-2776
Ad Production Coordinator: Kara Walby

Marketing & Circulation

Customer Service
Senior Director, Marketing Analytics:
Tricia Syed
Online Sales Development Director:
Amanda Phillips • 970-203-2806

Technology Division & Penton Marketing Services

Senior Vice President: Sanjay Mutha

Corporate

Chief Executive Officer:
David Kieselstein
Chief Financial Officer/Executive Vice
President: Nicola Allais



List Rentals

MeritDirect
333 Westchester Avenue,
White Plains, NY 10604

Reprints

Reprint Sales:
Wright's Media • 877-652-5295

Windows IT Pro, November 2012, Issue No. 219,
ISSN 1552-3136. *Windows IT Pro* is published monthly
by Penton Media, Inc. Copyright ©2012 Penton Media,
Inc. All rights reserved. No part of this publication may be
reproduced or distributed in any way without the written
consent of Penton Media, Inc.

Windows IT Pro, 748 Whalers Way, Fort Collins, CO 80525,
800-621-1544 or 970-663-4700. Customer Service:
800-793-5697.

We welcome your comments and suggestions about the
content of *Windows IT Pro*. We reserve the right to edit all
submissions. Letters should include your name and address.
Please direct all letters to letters@windowsitpro.com. IT pros
interested in writing for *Windows IT Pro* can submit articles
to articles@windowsitpro.com.

Program Code: Unless otherwise noted, all programming
code in this issue is ©2012, Penton Media, Inc., all rights
reserved. These programs may not be reproduced or
distributed in any form without permission in writing from
the publisher. It is the reader's responsibility to ensure
procedures and techniques used from this publication are
accurate and appropriate for the user's installation. No
warranty is implied or expressed.

Windows®, Windows Vista®, and Windows Server®
are trademarks or registered trademarks of Microsoft
Corporation in the United States and/or other countries
and are used by Penton Media, Inc., under license from
owner. *Windows IT Pro* is an independent publication
not affiliated with Microsoft Corporation. Microsoft
Corporation is not responsible in any way for the editorial
policy or other contents of the publication.

Windows IT Pro

IT Pros Are from Mars, Marketing Pros Are from Venus

What IT pros want to hear from technology marketers



**Sean
Deuby**

is technical director for
Windows IT Pro and *SQL
Server Pro* and former
technical lead of Intel's core
directory services team. He's
been a directory services
MVP since 2004.

Email



Twitter



In the midst of a wave of Microsoft marketing campaigns triggered by the company's major product announcements, I'd like to perform a public service. Let me address an important area that's widely misunderstood but rarely talked about: Bridging the gap between IT professionals and the marketing professionals who target them. It's a very real gap, and I'm convinced that one side doesn't realize how big this gap really is. Yes, I'm looking at you, marketers.

I spent many years working in an environment that had no interaction with vendors. (This is easy when you have no budget to purchase software.) Recently I've spent much more time around sales and marketing projects and people. I've sat in many meetings as "the technical guy"—or as I sometimes like to describe it to the 70 percent of marketers or marketing agencies that have never even met one, "Exhibit A: IT professional."

IT professionals and marketers are fundamentally different personality types. Unlike the human beings that marketers focus on, IT pros are usually more comfortable with technology than with groups of people. That's why they're in a field that deals with logical behavior and problem-solving instead of reaching out to irrational and unpredictable humans. IT pros favor dealing with predictability (things), and a marketer's job is to work with people—inherently

unpredictable. So let me give you a short list of what IT pros do and don't care about when shopping for technology.

- IT pros favor humor like Dilbert and [XKCD](#) (which has a second layer of humor in its alternate text when you hover over the cartoon) that has technical in-jokes.
- IT pros want technical information. IT pros tend to look down on anyone who isn't technical or in a technical profession. It's not pretty, but it's true. Marketing falls squarely into this category. One aspect of this: I'll bet easily half of the IT pros you ask won't even know the difference between sales and marketing, because it's not something they care about. If you can't stand toe-to-toe with an IT pro and talk at the 300 or 400 level he or she expects, be prepared for the cold shoulder.
- IT pros want just the facts. They have a finely tuned [band-pass filter](#) (I was an EE in a previous lifetime) that filters out BS on the low end and grand, overarching marketing themes on the high end. They don't even hear it. On a product landing page with a lot of text in the center column, I can guarantee you that most don't ever read it because they know it will talk about "redefining the server category," "transforming your IT operations," "[delivering a whole new level of business value](#)," and other statements that don't mean anything. Instead, they search on the outside edges of the page for "More Information" and "Technical Resources" so they can learn the technical features and how they might apply them. Even product datasheets are a bit flowery for them.
- IT pros do not want flashy presentations. Flashier is *not* better. IT pros don't mind some slickness. But beyond a point, high production values will work against you, not for you. IT pros respect a real person who's been in their shoes and who can help them learn something.
- IT pros want messages that are short and to the point. IT pros have what I term *technology-induced ADD*. They must absorb so much technical information on a daily basis, from multiple and

Unlike the human beings that marketers focus on, IT pros are usually more comfortable with technology than with groups of people.

increasing data sources, they have a very short attention span before they move on.

- IT pros are very logical and detail oriented. This mindset tends to be reinforced when, in my experience, a false move can lock out 30,000 users. Unlike practically anyone else, IT pros read manuals. For example, most of them know their way around their car's owner's manual. They actually keep owner's manuals for their personal electronics. Marketers should have a contingency plan for the detailed questions.
- “Just the facts” applies to video as well: Keep it short and sweet with high value (to them) content. Because IT pros are experienced manual skimmers, they're less likely to watch long videos because they're forced to follow at the presenter's pace rather than their own high-speed skimming pace.

I hope my rant helps the marketers out there understand their IT pro audience better. But you IT professionals aren't off the hook. Just because someone's in marketing—especially technical marketing—doesn't mean they aren't worthy of your time and respect. They might just have information you need to get the job done. They simply have a different skill set than yours. And they're probably more fun at parties!

What do you care about when evaluating technology and products? [Drop me a note](#) to share your point of view. ■

InstantDoc ID 144366

Letters

Generating Random Passwords

I want to thank Bill Stewart for the helpful PowerShell script that he provides in “[Generating Random Passwords in PowerShell](#)” (September 2012). Sometimes I find myself staring at the screen thinking of a completely random password. Great job!

—Bob Mitchell



Recommended Utility for Windows 8

I know a lot of users will complain about the lack of a Start button on the desktop in [Windows 8](#). But since installing Windows 8 and trying to like it, I finally gave in and installed a cool little utility called [Classic Shell](#). For me, Classic Shell has made the Windows 8 experience a whole lot better. I still use the Start screen but just not unnecessarily. I think your

readers will love it. I’ve been following Paul Thurrott from the start! Keep up the good work.

—Ed Jones

Get Your Windows 8 On

I read Paul Thurrott’s article “[Enterprises: Now’s the Time to Get Your Windows 8 On!](#)” and my response is that the Windows 8 “Metro” interface is not user-friendly for business enterprises. Up until Windows 8, I always enjoyed learning the new Windows experience of a new version. I’m a diehard Windows XP user, but I know I’m going to refresh my network with Windows 7 because it’s a great successor to XP. My users have already voiced their positive experiences about Windows 7. As for mobile devices, Windows 8 will probably turn out

Send Your Comments

Windows IT Pro welcomes feedback about the magazine. Send comments, and include your full name, email address, and daytime phone number. We edit all letters and replies for style, length, and clarity.



Comments

OK, but the “Metro” UI isn’t for laptops and desktops in the corporate environment.

Just for the fun of it, I installed Windows 8 on a system to test users’ reactions. Guess what? They all hate the interface. Not only did I get a negative response from everyone, but I believe Microsoft is making a huge mistake by not giving all of its consumer and business enterprise customers an option to use the Aero interface and Start menu. Worse is the fact that Microsoft is using the “Metro” interface on [Windows Server 2012](#), as mentioned in Michael Otey’s “[Is Microsoft Trying to Kill Windows Server?](#)” It’s not a good move on Microsoft’s part. The company is already alienating a large technical base in my area. I belong to a local technical group, and we’ve all agreed that we won’t purchase or install Windows 8 on any client system. ■

—Dave Scaletto

InstantDoc ID 144313

Hands-On Time with Windows Phone 8 and Why Windows RT Is Important

As we barrel toward the holiday season, questions continue to swirl around two mysterious Microsoft platforms that the software giant hopes will compete with emerging market leaders in the smartphone and media tablet markets. Yep, I'm talking about Windows Phone 8 and Windows RT, respectively. Although Microsoft remains mum on these topics, I'm ready to talk. Thankfully, I've actually had some hands-on time with each.

Windows Phone 8

My Windows Phone 8 experience, alas, has been virtual, courtesy of a leaked version of the Windows Phone 8 SDK that Microsoft, oddly, is keeping hidden from run-of-the-mill developers. The plan, apparently, is to seed the SDK early on to trusted or known developers only, to ensure that the quality of new apps is high on launch day. That day, incidentally, is October 29, 2012, another item of interest about which Microsoft, again, is being too quiet.

So what do we see in Windows Phone 8? As you probably know, this release is a huge architectural change. Microsoft is moving from the Windows CE core used in previous Windows Phone versions and is adopting the same core technologies (kernel and much more) used in Windows 8. This won't result in app compatibility—indeed, Windows Phone's WinPRT (the current name used for the phone-specific version of the Windows 8 Runtime) APIs diverge in some important ways from what developers will see in [Windows 8's](#) WinRT APIs—but it will have important ramifications for the scalability, capability, and performance of Microsoft's smartphone platform. (The dark side



**Paul
Thurrott**

is senior technical analyst for *Windows IT Pro*. He writes the SuperSite for Windows, a weekly editorial for Windows IT Pro UPDATE, and a daily Windows news and information newsletter called WinInfo Daily UPDATE.



Email



Twitter



Website

to this change is that Windows Phone 8 won't be offered as an update to existing devices. It requires, among other things, the multi-core processors found only in new hardware.)

What's most interesting to you, perhaps, is that while Windows Phone 8 is a sea change technically, it's really just an evolutionary update from a user-experience perspective. So if you're comfortable with Windows Phone 7.x, there's no jarring change moving up to Windows Phone 8, as there is, say, when you move from Windows 7 to Windows 8. Instead, you'll find several useful new features, none of which could be called earth shattering.

Since Microsoft is so keen to keep some of the new end-user features secret, I won't reveal everything I've seen. But some highlights include a new metered broadband connection feature called Data Aware that helps you track your data usage, a parental controls feature called Kid's Corner, a personal recommendation service for Local Scout, apps, games, music, Internet Explorer 10, a Wallet hub, a Nokia-infused Maps app, and of course a new Office 2013-based hub.

Only three hardware makers will be selling Windows Phone 8 handsets outside of China this fall, but they're all heavyweights: Nokia, HTC, and Samsung. Nokia and Samsung have already announced the devices they'll be selling, and HTC will probably have made its own announcement by the time you read this.

So far, I'm not sure what to think of these devices. Nokia will replace its flagship Lumia 900 with a 920 model that features an even bigger screen (with a correspondingly larger and heavier body), but the big change there could be the camera, which promises a PureView camera experience for superior picture taking. Its midlevel Lumia 820 model essentially combines the best aspects of today's Lumia 800 and 710 products, with a mid-sized screen, expandable storage, and replaceable back covers for a customized look.

Samsung's single Windows Phone 8 entry, the ATIV S, is a follow-up to last year's excellent Focus S, and if you're familiar with the firm's Android-based Galaxy S III, you get the idea: It will feature a huge 4.8"

HD screen, a 1.5GHz dual-core processor, 1GB of RAM, an 8 megapixel rear camera, and expandable storage. It doesn't look like much, but given Samsung's track record, this could be one to watch.

HTC hasn't announced its Windows Phone entry. However, all the rumors point to a Windows Phone version of its excellent Android-based One X handset, so that's probably going to be a decent device as well.

An embarrassment of riches? Perhaps. But Windows Phone 8 is entering a crowded market, one that just became a lot more crowded thanks to the entry of Apple's new iPhone 5. Nothing about the new iPhone is earth shattering; indeed, its two best features—large screen and LTE networking support—have been available on Windows Phone and Android handsets since last year.

But Apple's spend-happy fans have proven open to buying up even the most lackluster of upgrades in the past—cough, iPhone 4S, cough—so this should be no different. It's fair to think that Apple will sell millions of these things immediately.

On the Android side, Motorola Mobility's recent RAZR announcement didn't impress all that much, but the current bestsellers, the Samsung Galaxy S III and HTC One X, are already strong enough to take on the iPhone 5—and of course the coming Windows Phone 8 handsets—head to head. But what makes Android so successful is the same thing that worked for Microsoft and its Windows PC ecosystem: There's such a wide diversity of devices out there, available on all mobile carriers, that Android has simply become the obvious choice in every market segment. It's a tough opponent to overcome, and all the pieces are in place for Android to continue its dominance of the smartphone market.

Windows 8 and Windows RT

Things are quite different on the tablet side. There, the dominance is reversed, with Apple and its iPad controlling 65 percent of the market and various Android makers sopping up most of the rest. In

this important emerging market, Microsoft is offering a two-pronged attack featuring Windows 8 and Windows RT.

Windows 8 is pretty well understood by this point. It runs on PCs, and for this generation of hardware, we're going to see a ton of new device types, not just traditional laptops, ultrabooks, and desktops, but also a lot of multi-touch-enabled screens, and hybrid mobile devices such as slate PCs, tablet PCs, convertible PCs, and more.

Windows RT is a different animal entirely. Although Microsoft has been careful to keep the wraps on its ARM-based variant of Windows 8, I was able to spend a few days with a Qualcomm Liquid reference design tablet with dock, and I think it's fair to say I know what's going on here now. It's simple: Windows RT is nothing less than a complete rethinking of what Windows can and will be in the future. It's the future, in the same way that Windows NT was the future in the mid-1990s.

Put another way, Windows RT is the device-based version of Windows 8. (Or, Windows 8 is the PC-based version of Windows RT.) There are differences, few of which are subtle. Windows RT lacks some Windows 8 desktop features—Windows Media Player and Storage Spaces—and can't run any non-bundled Windows desktop applications, such as Adobe Photoshop and Visual Studio. But in return, you get a clean new version of Windows that lacks legacy deadwood and the attendant security and reliability issues, has killer battery life, and ships on devices that are silent, thin, and light.

If you choose to allow Windows RT into your environment, however, you don't need to give up backward compatibility. It's no coincidence that Microsoft has spent the past several years honing its centralized app deployment (RemoteApp) and data-center-based Windows environment (VDI) solutions. Both work just great with Windows RT. In fact, I mentioned Photoshop for a reason: My Windows RT test included a RemoteApp version of Photoshop, as well as a full VDI-based Windows desktop that could run a 3D CAD-CAM application and play full-screen HD video simultaneously.

Various hardware makers will be selling Windows RT devices and Windows 8 PCs of all kinds come late 2012. But none, perhaps, are as eagerly awaited as the Microsoft Surface devices. Microsoft will sell two models, one based on Windows RT (which will go on sale October 26, 2012) and one based on Windows 8 (which won't see the light of day for another 90 days, or roughly until February 1, 2013).

To date, no one outside of Microsoft has had any hands-on time with actual working Surface devices—despite some rather sad and quickly debunked claims to the contrary—but this has only heightened the aura and excitement around these devices. That's quite a trick, given Microsoft's normally lackluster marketing efforts. But as good as these efforts look, there are just too many questions, not the least of which is price: How much will Microsoft sell these things for? We just don't know. Not yet.

Here, Microsoft is also entering a crowded market. Although Surface RT and other Windows RT devices won't compete with most of these products head-to-head, it's fair to say that Windows RT's biggest competitors are the iPad and various Android-based tablets.

Consider this: While Apple's iPad doesn't compete with a traditional PC in any feature-by-feature comparison, it does compete with a PC in the sense that some people can get real work done with the iPad and don't need the complexity or full functionality of a PC. In the same way, even simpler media tablets, such as the Google Nexus 7 or the 7" Amazon Kindle Fire/Fire HD, might not line up toe-to-toe with the iPad. But they do compete with the iPad because every Nexus 7 or Kindle Fire sale is a non-sale for Apple. The fate of all these devices is in some way interconnected.

Where does Windows RT fit into this comparison? Devices based on Windows RT occupy a previously unknown middle ground between PC and iPad, and these devices will thus compete both with true Windows 8 PCs (and slates and other form factors) and with the iPad and full-sized Android tablets. But devices based on Windows RT don't compete directly with smaller, simpler media tablets such as the

Nexus 7 or Kindle Fire, because those devices are purely for media and content consumption.

So, looking at competitors for the market the Surface and other Windows RT devices are aimed at, we see . . . the iPad 3. Yes, there are 10"-ish Android tablets out there too, but they either don't sell well enough to worry about or, in the case of the new 8.9" Amazon Kindle Fire HD, they're just too purely media focused, lacking any productivity accoutrements at all, to be included in this comparison.

The iPad 3 is strong competition. It has an established and superior ecosystem of content, apps, games, and accessories. Consumers love it and the Apple brand. And yet, the iPad has lost share every year since its launch, even though its sales have only improved. Those Android tablets—especially Amazon's devices—have already had their effect. Perhaps Windows RT can make some serious inroads here.

Looking Forward

I'd like to leave you with two final thoughts. First, Apple is expected to unveil a so-called iPad Mini that will compete with the Nexus 7 and Kindle Fire, providing a true media tablet alternative. This is important for Apple, which needs to fend off its voracious competitors. It's not so hugely important to you.

Second, the dark horse for Microsoft—go figure—is device management. Whereas all of the devices and platforms that I mentioned can be managed through Exchange ActiveSync (EAS), Microsoft's platforms will offer better management capabilities. Windows 8, of course, can be managed through Active Directory (AD) and Group Policy. Windows Phone 8 and Windows RT will provide additional management functionality, surfaced through System Center 2012 SP1 and the next version of Windows Intune. If you've grudgingly accepted BYOD (Bring Your Own Device) in your own environment, these capabilities could put Windows Phone 8 and Windows RT over the top. It's something you need to look into. ■

InstantDoc ID 144260

Automating PowerShell Reports, Part 2

Using Windows Server's "hidden" SMTP server to email anonymously ... and safely

In “Automating PowerShell Reports, Part 1,” I introduced one of my favorite PowerShell cmdlets, *send-mailmessage*. It’s a flexible tool that’s essentially a command-line SMTP client, allowing you to send Internet email from the command line. I introduced it because I’ll soon be showing you how to use Active Directory’s (AD’s) cmdlets to create AD reports, schedule your servers to run those reports in the middle of the night, and automatically email you those reports. But you won’t be able to do all that unless you enlist the aid of a great power tool, Windows Server’s SMTP Server module.

Say you have a routine that runs every night at 3:00 A.M., generating a report called `C:\reports\today.html`, and you want that report automatically emailed to you at the email account `admin@bigfirm.com`. As I demonstrated last month, that command might look like

```
send-mailmessage -from automaticjobs@bigfirm.com
    -to admin@bigfirm.com -subject "Daily Active Directory Report"
    -body "C:\reports\today.html"
    -smtpserver ourmailserver.bigfirm.com -bodyashtml
```

Try running that from a PowerShell command prompt, and it’ll probably fail because well-run email servers don’t send mail except from authenticated users, which means (in practice) that every time you run *send-mailmessage*, PowerShell is going to ask you for credentials.



Mark Minasi

is a senior contributing editor for *Windows IT Pro*, an MCSE, and the author of 30 books, including *Mastering Windows Server 2008 R2* (Sybex). He writes and speaks around the world about Windows networking.



Email



Twitter



Website

Unfortunately, you won't be around to provide those credentials, because the routine is running in the background at 3:00 A.M. on some server. What to do? My suggestion is that you don't Google this question; many pages direct you to download and install a random free SMTP server that might or might not be securable. The Internet definitely doesn't need another "open relay," an email server just waiting for some spammer to use it to blight several million mailboxes with junk, and who would want to live with that kind of karma anyway?

Instead, install the SMTP Server service that has shipped with Windows Server for ages, and then configure it so that it doesn't need authentication—but also isn't an open relay. The SMTP Server service used to be a marquee feature of IIS through IIS 6.0, and it has always been a favorite of mine because it's flexible, easy to configure, and programmable. You can actually build things such as spam filters quite easily with VBScript! With Windows Server 2008 and IIS 7.0, however, Microsoft quietly and inexplicably removed it. (Was the reason to boost Exchange Server sales? That's a scary thought!) The company did, however, include the old IIS 6.0 (Windows Server 2003) SMTP module, although installing it from the IIS 7.0 administrative GUI has always eluded me.

With PowerShell, however, it's a snap:

1. Start PowerShell, and type

```
import-module servermanager  
add-windowsfeature smtp-server  
add-windowsfeature rsat-smtp
```

That's a different *import-module* than the one you normally use, but it loads a few cmdlets that let you turn server roles and features on and off. (*Get-windowsfeature* lets you see what's available and what's currently enabled, and *remove-windowsfeature* lets you disable roles, role services, and features.)

2. Now that you have the SMTP server installed, you can secure it from the old Server 2003 IIS Manager. Click Start, All Programs, Administrative Tools, Internet Information Services (IIS) 6.0 Manager to bring that up.
3. In the resulting IIS Manager window, you'll see an icon of a server with your server's name and *(local computer)* next to it, with a plus sign alongside it. Click that plus sign, and you'll see an icon that looks like a brownish envelope with speed lines emanating from it. The icon is labeled *[SMTP Virtual Server #1]*.
4. Right-click that, and choose Properties to access the multi-tabbed *[SMTP Virtual Server #1] Properties* page. Click the Access tab.
5. On that tab, click Relay to open the Relay Restrictions dialog box.
6. In the Relay Restrictions dialog box, clear the check box labeled *Allow all computers which successfully authenticate to relay, regardless of the list above*. In that same dialog box, ensure that the *Only the list below* radio button is selected (it's the default, so it should already be), and then tell the server to accept requests only from local processes by clicking Add to open a dialog box labeled Computer.
7. Under *Add one of the following to the list*, select the *Single computer* radio button. In the *IP address* field, fill in 127.0.0.1.
8. Click OK to close the dialog box, close the Relay Restrictions dialog box, and close the *[SMTP Virtual Server #1] Properties* dialog box.

I'm hoping that, if I hadn't yet demonstrated why I like the command line for administration, this bit of configuration makes the point! Now you can safely run a *send-mailmessage* command on that server, being sure to specify *-smtpserver localhost*, and you won't be asked to authenticate. Now we're ready to get that report rolling, and we'll do that next month. ■

InstantDoc ID 144192

FAQs About Microsoft Virtual Academy

Keep your Microsoft skills up-to-date with free online courses



**Michael
Otey**

is senior technical director
for *Windows IT Pro* and
SQL Server Pro and author of
*Microsoft SQL Server 2008 High
Availability with Clustering &
Database Mirroring*
(McGraw-Hill).

Email



Keeping up with technology is often one of the toughest parts about being in IT. Technology is constantly changing and evolving, and you need to stay on top of the latest developments to keep your skill levels current. Although a number of different training venues are available, one of the least well-known, but also one of the most easily accessible, is [Microsoft Virtual Academy](#) (MVA), an online portal for IT pros to learn about Microsoft's cloud-based technologies. In this column you'll get answers to the top 10 FAQs about MVA.

① What does MVA cost?

Believe it or not, Microsoft Virtual Academy is completely free. Training is one of those areas where Microsoft really provides extra value to IT professionals. Naturally, the courses focus on Microsoft products, but that's what you'd expect from Microsoft training. The only requirement is that you have a Windows Live ID.

② If the courses are free, are they valuable and up-to-date?

MVA covers topics in a wide range of Microsoft technologies, and there are different levels of courses. Some examples of the current courses include "Windows Server 2012: First Look," "System Center 2012: Orchestrator & Service Manager," "Microsoft Licensing Fundamentals," and "System Center 2012: Virtual Machine Manager." Although I noticed the material for the "Windows Server 2012: First Look" course was based on the beta and didn't reflect the latest updates, it still provides a good overview of the new features.

③ How many people have taken courses from MVA?

For such a little-known offering, MVA has a surprising number of participants. Although these numbers are always changing, to date Microsoft states that 768,500 students have registered for MVA courses and that 826,200 self-assessment exams have been passed.

④ How are the courses organized?

The courses are typically organized into multiple modules with a self-assessment exam at the end of each module. For instance, the “Windows Server 2012: First Look” course has the following modules: “Windows Server 2012 Overview,” “Beyond Virtualization,” “The Power of Many Servers,” “Modern Workstyle Enabled,” and “Every App, Any Cloud.” Successful completion of each module accumulates a certain number of track completion points.

⑤ Do the results count toward certifications?

No. However, the training can certainly help with taking other certifications. The MVA program tracks accumulated points for successfully completing different modules. You see a dashboard when you sign in that tracks your point progress for Bronze (0-99), Silver (100-499), Gold (500-2,999), and Platinum (3,000+) levels.

⑥ Are you limited in how many courses you can take at any given time?

No. You can take any number of MVA courses at any time, depending on availability. You can also retake each of the self-assessment exams as many times as you want.

⑦ What are the courses like?

The courses consist of multiple modules. Each module typically contains a collection of different resources. Some MVA modules are video presentations that you can stream over the web, and others are Word or PDF documents that you can read online or download and review

Training is one of those areas where Microsoft really provides extra value to IT professionals.

offline. In case you're wondering, the MVA modules don't let you skip right to the self-assessment exam—you must review some of the study materials first.

⑧ **How many tracks and courses are available?**

Microsoft is continually adding and updating the content of the MVA courses. At the time of this writing, there are 15 major tracks: Business Intelligence, Licensing, Office 365, Private Cloud, Public Cloud, Security, SQL Azure, SQL Server, System Center 2012, VDI, Virtualization, Windows Azure, Windows Client, Windows Server 2008 R2, and Windows Server 2012. Each track has at least two courses, and many have more. There are 39 total courses.

⑨ **Where do you register for MVA training?**

You can find [Microsoft Virtual Academy](#) online.

⑩ **How do you get started taking MVA courses?**

You get started by signing in to MVA with your Windows Live ID and then filling out the registration page. At that point, a confirmation email is sent to the email account that you supplied during registration. Clicking on the confirmation link in the email opens your new MVA account page. From there you can begin selecting courses. ■

InstantDoc ID 143921

Building Your Identity Bridge to the Cloud

Securely connect your on-premises identity with cloud services

Most identity professionals are highly focused on their particular area of responsibility and don't have a lot of time to broaden their knowledge to other areas that are related to—but don't directly affect—their day-to-day jobs. This makes it challenging to have an awareness (let alone an understanding) of the different types of identity services that have sprung up in recent years. As the maxim goes, “You don't know what you don't know.”

In May, Gartner analyst Mark Diodati published a report entitled “[Identity Bridges: Uniting Users and Applications Across the Hybrid Cloud](#).” This report summarizes the evolution of a product segment that I've been writing about in this column since I joined *Windows IT Pro*. Diodati calls this product an *identity bridge*, which is an appropriate name because it bridges the traditional on-premises Identity and Access Management (IAM) system such as Active Directory (AD) to cloud-based services that need these identities. All working identity professionals need to become conversant with the capabilities encompassed by identity bridges, because they will be expected to understand and perhaps recommend what their company needs in order to implement its own bridges to the cloud.

Incidentally, Diodati might not have intended it, but I can't resist mentioning that the term identity bridge also nicely parallels Norse mythology. (No, I have never attended Comic-Con, but I did grow up watching *Thor* cartoons after school.) As you can see in Figure 1, [Bifrost](#) is also a bridge: the rainbow bridge to the Norse god's home



Sean Deuby

is technical director for *Windows IT Pro* and *SQL Server Pro* and former technical lead of Intel's core directory services team. He's been a directory services MVP since 2004.



Email



Twitter

Figure 1
Bifrost, the Original
Identity Bridge



of [Asgard](#). Bifrost connects the world we're familiar with (mortals, or on-premises identity) with the unfamiliar, ethereal world in the clouds (immortals, or cloud-based services). You could stretch the Norse analogy even further to say that the god [Heimdallr](#) is the equivalent of the identity router component of many identity bridges, guarding the bridge against transgressions. I'm surprised that [Nordic Edge](#), a federation software company that [Intel](#) purchased to jumpstart its entry into this market, never picked up on this analogy. Now you can go watch *Thor* and say you're doing identity research!

Diodati's report nicely summarizes the many different types of products and capabilities that might be required to provide seamless and secure identity to cloud services, from the well known and mature to the still emerging. Although vendors vary in their products' competencies, some of the capabilities that identity bridges can provide are as follows:

- Federation for single sign-on (SSO) authentication (AuthN) to the cloud service. A key component of this capability is the ability to transform security tokens from a standard accepted in one realm (e.g., Kerberos tickets in an AD environment) to a standard accepted in another realm (e.g., SAML tokens in a web service environment or OAuth tokens in a mobile environment). For more information, see my article "[Ease Cloud Security Concerns with Federated Identity](#)."

- Directory synchronization. This is typically one way, from the identity provider (e.g., your company's AD implementation) to the service provider (the target cloud service) to ensure that changes made to the identity provider, such as disabling an account, are immediately replicated to the service provider. For more information, see my article "[Identity Predictions](#)."
- Just-in-Time (JIT) provisioning. A JIT provisioning capability ensures that an account is created at the service provider only when a user first attempts to access the service. Among other advantages, this means that a company isn't charged for user access to a service until it actually begins using it. Note, however, that JIT provisioning covers only the creation of the account; updates to and deletion of the account must be handled by another method.
- Authorization (AuthZ) services to determine who can access which services.
- Virtual Directory Services. VDS provides an aggregate view into what might be many separate enterprise identity stores. For more information, see my article "[The Rise of Virtual Directory Servers](#)."
- Password vaulting. Although identity federation is the strategic direction for providing secure authentication for cloud services, the reality today is that smaller Software as a Service (SaaS) providers aren't generally set up to support federation. Instead, they rely on inputting a user ID and password to authenticate users. Password vaulting stores a user's credentials in the identity service and replays them to the SaaS website as if the user was directly logging on. (Incidentally, this is how the [LastPass](#) browser add-on provides its autofill and autologin authentication functions.) I've seen the password-vaulting capability mentioned only as part of Identity as a Service (IDaaS) solutions.

Some identity bridges—for example, Microsoft's Active Directory Federation Service (AD FS)—don't have a SaaS component

([Windows Azure AD notwithstanding](#)), but many are tightly integrated with an [identity-management service](#) or are entirely service-based. Mobile device management, which can distribute authentication credentials to smartphones and tablets, has quickly become a standard capability.

Like many markets, the identity bridge market was pioneered by startups that offered unique, single-purpose products. As this market matures, the startups grow into larger companies with broader product portfolios and greater capabilities in their products, often combining two or three identity-related capabilities. These small companies are often acquired by larger players seeking to jump into this area or to quickly add capability to their existing products. This maturity also gives rise to a class of products that Diodati calls *super bridges*, which have a superset of services such as storage and network load balancing that go beyond identity services alone.

Reading Diodati's report is a great way to quickly understand the different aspects of these identity bridging technologies and who the key players are in the market. The market is evolving rapidly, so although most of the base technologies will still apply at this time next year, I expect that many of the players will have changed. And as if to underscore these changes—and his faith in the future expansion of this market segment—Diodati left [Gartner](#) in August to join [Ping Identity](#). Download the “[Identity Bridge + Identity-as-a-Service: Where will it take you?](#)” report from Ping Identity's website. I'll continue to pay close attention to this area and keep you aware of interesting developments and information you need to know. ■

InstantDoc ID 144308

Reader to Reader

Disable Video ActiveX Controls in Internet Explorer on the Fly

One of my customers complained to me about the “slow Internet” on her Windows XP SP3 machine, which was running Internet Explorer (IE) 8.0. Internet speed tests showed no problem, so I looked into the kinds of websites she was visiting and noticed all the Adobe Flash content. These days, it’s hard to find websites that don’t have video ads.

I realized that the video ads were likely the cause of the problem, so I decided to disable videos in the browser. My first attempt was to simply clear IE’s *Play animations in web pages* check box. (This check box is in the Multimedia section on the Advanced tab of the Internet Options page, which you access through the Tools menu.) When that didn’t work as expected, I did a little research and found that this setting affects only animated GIF files.

After doing even more research, I discovered that there are a few fast and easy tweaks you can make to improve users’ Internet experience and productivity on older, slower systems (e.g., Windows Vista, XP) that use IE 8.0 or IE 9.0. These tweaks will also help with network bandwidth utilization.

Here’s what you need to do:

1. Open IE and go to a website that has Flash video (e.g., YouTube.com). Then, go to a website that uses Microsoft Silverlight (e.g., Microsoft.com). By going to these websites, you’ll be loading the ActiveX controls for Flash and Silverlight.
2. Select *Manage Add-ons* on the Tools menu. Make sure that *Currently loaded add-ons* is selected in the Show drop-down



Bret A. Bennett

is the principal consultant of BRET A. BENNETT. He specializes in the installation, training, and support of Microsoft Dynamics GP, Passport Business Solutions, and Open Systems Traverse software, and he provides support for BI delivery and Microsoft servers.



Email

- list. In the right pane, find Shockwave Flash Object and click it. Use the Disable button to turn it off. Then find Microsoft Silverlight and click it. Use the Disable button to turn it off.
3. Close the *Manage Add-ons* dialog box.

With this technique, you can disable videos (and other add-ons) quickly and on the fly, without exiting IE. If you want to see active video in your IE session again, just go back to the *Manage Add-ons* dialog box, enable the ActiveX controls, and use the F5 key to refresh the web page.

If manually disabling the ActiveX controls isn't feasible, you can implement this technique with Group Policy. For information on how to do so, see the *Windows IT Pro* article "[Managing Windows Vista Group Policy Options](#)" (March 2007) and the Microsoft article "[How to manage Internet Explorer add-ons in Windows XP Service Pack 2.](#)" ■

InstantDoc ID 143876

FAQ

Answers to Your Questions

Q: What is Windows 8 File History?

A: Windows 8 File History is a new feature that protects data files on your computer in the Libraries location and on the desktop, in addition to your Favorites and Contacts. Once an hour (by default), it copies any changes to an alternative location.

In addition to just copying new data, a history of the data is maintained on the alternative location. You can view that data for previous points in time via the History button in Explorer. History shows the different versions of Libraries, folders, and individual files.

File History is enabled and managed with the File History Control Panel applet. On the Advanced Settings page, which Figure 1 shows, you can configure the amount of space to use as a cache (used when the target for File History, such as an external hard drive, isn't available), how often to save copies of files, and how long to keep saved



Jan De Clercq



William Lefkovics



John Savill

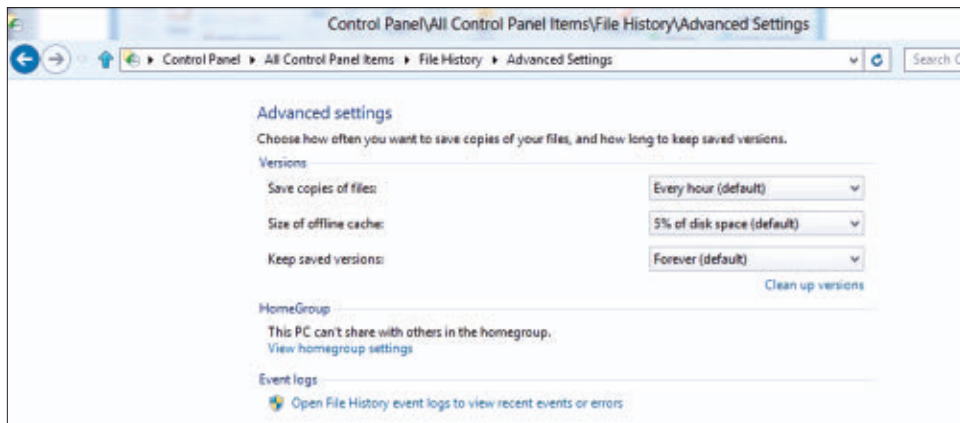
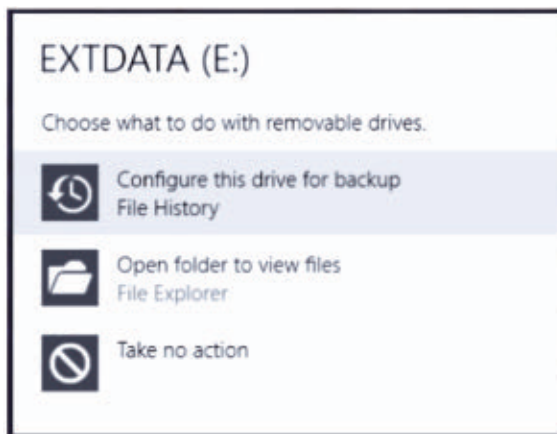


Figure 1
Windows 8 File History Advanced Settings

Figure 2
Windows 8 File
History Dialog Box



versions (the more history that's stored, the more disk space required).

The File History setup is also available when a new hard disk is attached to the system, at which point one of the options is to use the new drive for File History. Selecting this option launches the

File History configuration dialog box (see Figure 2).

File History can be disabled by going to Computer Configuration, Administrative Templates, Windows Components, File History, and selecting the Group Policy setting *Turn off File History*.

—John Savill

InstantDoc ID 143933

Q: How can I set up RMS-based protection for the documents users store in SharePoint?

A: You can use Windows Rights Management Services (RMS) to protect SharePoint documents in the two most recent releases of SharePoint: SharePoint Server 2010 and SharePoint Server 2007 both include RMS support. However, you should be aware of some restrictions and complexities if you plan to set up RMS with your SharePoint installation.

An important thing to know is that RMS can encrypt SharePoint documents and subject them to RMS access control restrictions only when they are downloaded from a SharePoint 2010 or SharePoint 2007 document library. RMS doesn't leave SharePoint documents encrypted while they're stored on the SharePoint server. This restriction exists so that SharePoint can index and scan the documents on a SharePoint

storage provider. RMS applies its restrictions to a document only right before it's downloaded to a client computer. Similarly, when an RMS-protected document is uploaded to a SharePoint site, RMS removes all protection from the document until a new download request is received.

SharePoint-RMS integration ensures that security restrictions are enforced even after a document has left a SharePoint server, which is something that can't be achieved using the standard SharePoint permissions. SharePoint-RMS integration also automatically enforces an organization's RMS document security policies. A SharePoint administrator can centrally define different RMS policies for the document libraries hosted on a SharePoint server. Therefore, individual users don't have to decide what protection they need to apply to documents they post in SharePoint libraries. RMS permissions are defined at the SharePoint document library level: Documents in a library automatically inherit the library's RMS permissions. This protection applies to both existing and new documents in the SharePoint library.

The RMS protection of SharePoint data is, just like the RMS protection that's bundled with Windows and Microsoft Office, possible only for certain file formats. Out of the box, it supports Word, Excel, PowerPoint, InfoPath, and XPS files. Extensions to apply RMS protection to other file formats (e.g., .pdf, .cad) can be added through special software from Microsoft partners such as TITUS and GigaTrust.

RMS support for SharePoint can be set up using either RMS SP2 or RMS 2.0, which is bundled with Windows Server 2008. Provided you already have a functioning RMS infrastructure, enabling RMS protection in SharePoint is relatively straightforward. The main configuration actions are

- enabling RMS support on the SharePoint server
- setting the actual RMS restrictions in the configuration of a given document library

You can enable RMS support in SharePoint by selecting either the *Use the default RMS server specified in Active Directory* or *Use this RMS server* option in the Information Rights Management section of the \SharePoint Central Administration\Operations configuration section.

To set RMS restrictions on a SharePoint document library, you must use the Information Rights Management section in the *Permissions and Management* configuration section of the document library. When you select the *Restrict permission to documents in this library on download* check box, you can further refine the RMS protection as follows:

- Allow users to print documents.
- Force users to verify their credentials every x number of days. This setting can be useful when someone who has access to RMS-protected confidential data leaves your organization; the individual will retain access to the data only for x days after his or her last successful authentication to an RMS server.
- Reject files that don't support Microsoft Information Rights Management (IRM). Selecting this option results in SharePoint rejecting the upload of document formats that don't support RMS.
- Remove RMS protection on a particular date. This setting is useful for publishing company financial results, for instance. After the quarterly results are published, the RMS protection policy on the quarterly results SharePoint library automatically changes—meaning that the RMS restrictions are removed.

For detailed guidance on how to set up SharePoint-RMS integration, see the Microsoft article [“Deploying Windows Rights Management Services with Microsoft Office SharePoint Server 2007 Step-By-Step Guide.”](#)

—Jan De Clercq
InstantDoc ID 144305

Q: Can you reclaim a task in Microsoft Outlook after you've assigned it?

A: One of the great collaborative features in Outlook is the ability to generate tasks and assign them to individuals. Tasks can be simple or complicated. They might be composed of many details that make them time-consuming to recreate. Therefore, it would be nice if you could reclaim a task that you've previously assigned that hasn't been completed, but Outlook currently doesn't give you that ability.

When a new task item is created, you can assign the responsibility of the task to another user. Simply select the Assign Task button in the Ribbon, as Figure 3 shows.

This function adds a To field to the Task form, as Figure 4 shows. After you've filled out the Task properties as desired, you then assign the task and send it, like a message, to the desired recipient.

Note that when you assign a task from an account using Microsoft Exchange Server, there's a check box, selected by default,

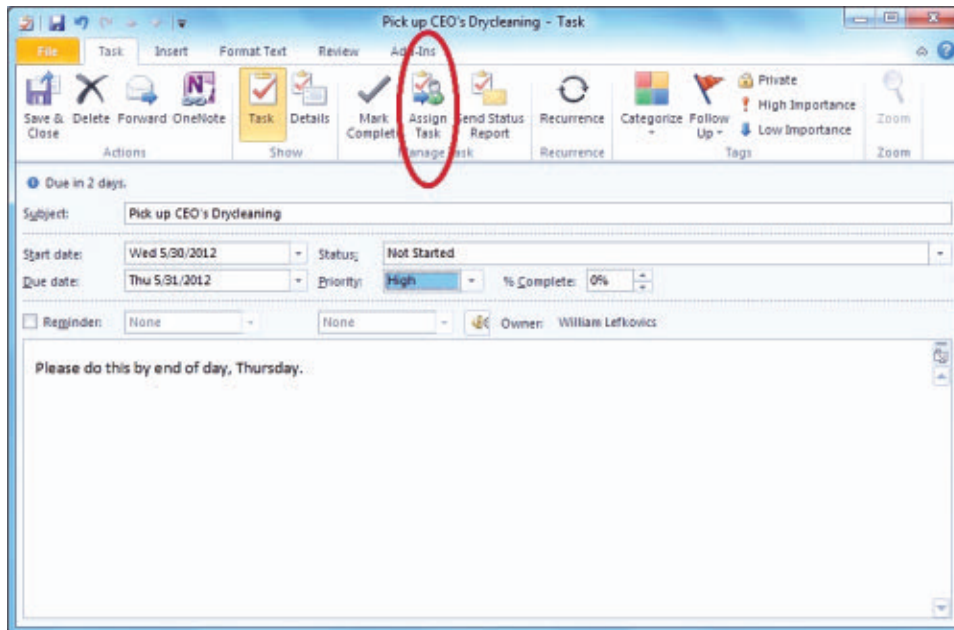
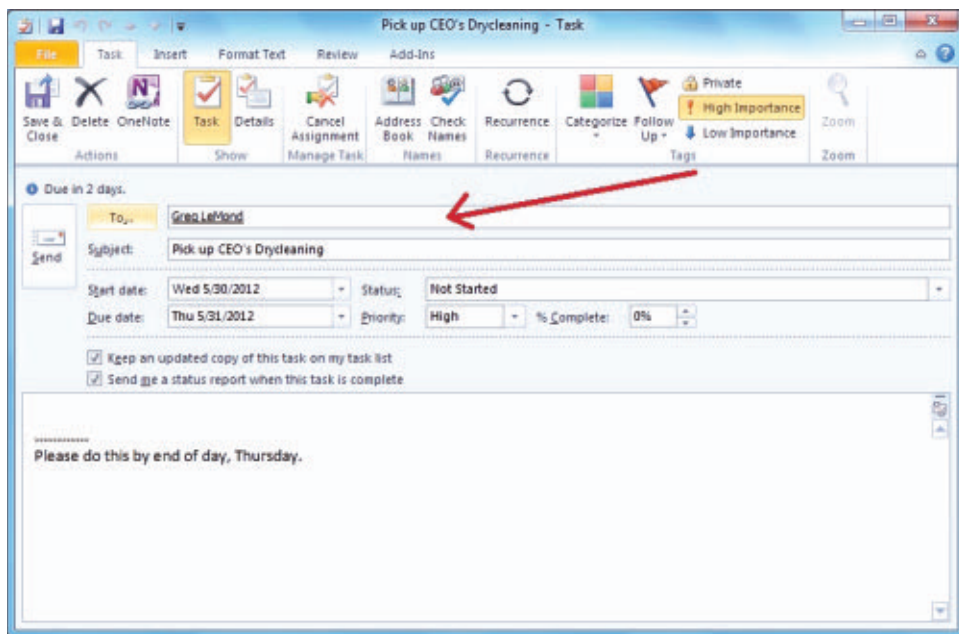


Figure 3
Assign Task Button
on a New Task in
Outlook

Figure 4
Assigning an Outlook
Task to a User



requesting that a copy of the updated task reside in your task list. Using this option can help you recreate the task if you need to reassign it later.

Incidentally, I had a user who clicked *Save & Close*, thinking that this action assigned the task. If you were adding a task to your own task list, you would click the *Save & Close* button, but you must click *Send* to actually send an assigned task to the specified recipient.

After the task is sent, you must wait for the recipient to accept or decline the task. At that point, you, the task assigner, will be notified by a return message to your Inbox. Assuming the task is accepted, the task is no longer yours. There's no way to recall the task or reassign it without action by the assigned recipient. However, the recipient can decline a task that's already been accepted, sending the task back to the sender, who can then reassign it. By the way, if the recipient is using Outlook Web App (OWA), he or she won't be able to accept or decline an assigned task. As you can see in Figure 5, this action requires the Outlook client.

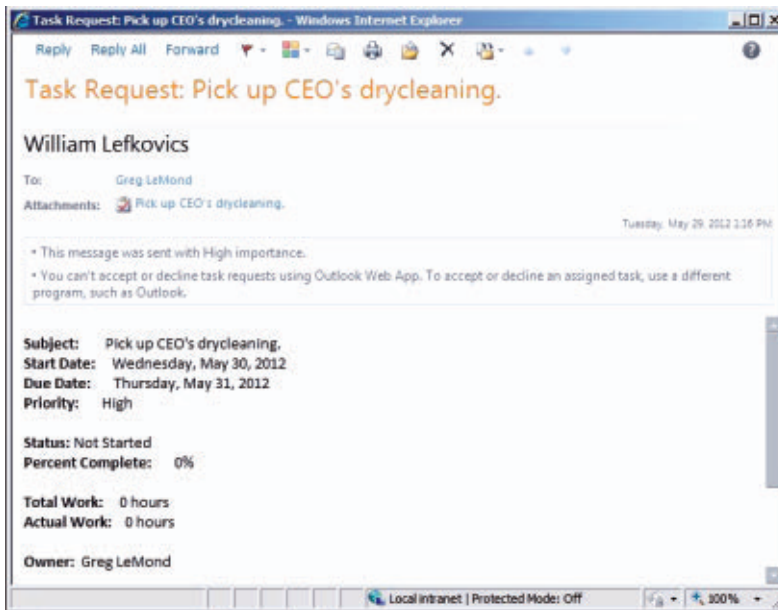


Figure 5
Outlook Task Request
Viewed in OWA

When a recipient of an assigned task declines said task, the task is returned to the sender, who can then select *Return to Tasks* to reassign the same task to another user. If a recipient has accepted a task but is no longer available to complete or decline the task (e.g., he or she has been let go from the company), you might not want to recreate the entire task. Assuming you left the check box selected that leaves a copy in your task list, you can right-click that assigned task and select Copy, then paste it in the white area below your task list. In Outlook 2007, you can highlight the task, select Edit, *Copy to Folder*, then choose the same Tasks folder to place the copied task. This operation generates a new, unassigned task with all the details intact. You can then delete the assigned task that you're unable to reclaim from the initial recipient. Alternatively, you or an authorized administrator can log on to the assigned recipient's mailbox, if it still exists, and decline the task to free it up for reassignment to another recipient. ■

—William Lefkovic
InstantDoc ID 144081

Microsoft Releases Windows Server 2012

Improvements in storage, virtualization, and management are worth a look

Windows Server 2012, arguably the most significant server release Microsoft has ever offered, became available for evaluation and purchase to customers around the world on September 4, 2012. Server 2012 offers a simplified licensing model that includes all features of the OS in all editions of Server. You'll find improved management capabilities in Server Manager and PowerShell. Storage improvements are numerous, and Hyper-V enhancements include scalability, live migration upgrades, and storage live migration capabilities. *Windows IT Pro* brings you ongoing coverage of Server 2012, with in-depth treatment of significant features, breaking news, and analysis. Visit our [Windows Server 2012 page](#) for the latest news and technical features. ■

InstantDoc ID 143935

Top 10 Windows Server 2012 FAQs

- 1 What is Offloaded Data Transfer in Windows Server 2012?
- 2 After I reinstalled Windows Server 2012, my Storage Spaces are no longer writable or automatically attached—what can I do?
- 3 Can I upgrade a Windows Server 2008 or Windows Server 2008 R2 Server Core installation to Windows Server 2012 with a GUI directly?
- 4 What Windows PowerShell cmdlet adds a VHD to a virtual machine in Windows Server 2012?
- 5 Why, when I enable .NET Framework 3.5 on Windows 8 and Windows Server 2012, does it connect to the Internet and pull down files?
- 6 What is the Windows Server 2012 NUMA Spanning option, and should it be enabled or disabled?
- 7 Does SMB Transparent Failover in Windows Server 2012 require ReFS?
- 8 Does Windows Server 2012 Essentials or Foundation include the Hyper-V role?
- 9 Why would you want to make a Windows Server 2012 Scale-Out File Server cluster access a SAN via the file servers using SMB 3.0?
- 10 How do I enable and view the Windows Server 2012 Hyper-V metric information?

Windows Server 2012 Articles

- ▶ [Introducing Windows Server 2012](#)
- ▶ [New Features in Windows Server 2012 Server Manager 2012](#)
- ▶ [Windows Server 2012 Sprints Through the Finish Line](#)
- ▶ [Getting Around in Windows Server 2012, Part 2: Server Manager](#)
- ▶ [Get Ready for Windows Server 2012 Hyper-V](#)
- ▶ [Cloning Virtual Domain Controllers in Windows Server 2012](#)
- ▶ [Video: Getting Around in Windows Server 2012 Server Manager](#)
- ▶ [Windows Server 2012 Essentials: Connect Client PCs without Using a Domain](#)
- ▶ [New Ways to Enable High Availability for File Shares](#)
- ▶ [Microsoft Releases Windows Server 2012 to Manufacturing](#)
- ▶ [Top 10 Windows Server 2012 Storage Enhancements](#)
- ▶ [Is Microsoft Trying to Kill Windows Server?](#)
- ▶ [Getting Around in Windows Server 2012, Part 1](#)
- ▶ [Shared-Nothing VM Live Migration with Windows Server 2012 Hyper-V](#)
- ▶ [Windows Server 2012 Simplifies Active Directory Upgrades and Deployments](#)
- ▶ [Windows Server 2012 Storage Spaces](#)
- ▶ [Video: Windows Server 2012 Storage Spaces Demo](#)
- ▶ [How Windows Server 2012 Improves Active Directory Disaster Recovery](#)
- ▶ [Introducing a Simpler Windows Server](#)
- ▶ [Windows Server 2012 Will Have Feature Parity Across All Editions](#)
- ▶ [Windows Server 2012 Is Good News for IT](#)
- ▶ [Top 10 New Features in Windows Server 2012](#)

Sponsored by ▼

EMC²

www.windowsitpro.com/windows-server-2012

Secure the Private Cloud

5 steps that your organization can take now



**John
Howie**

is the COO of the Cloud
Security Alliance.

Email



Twitter



LinkedIn



The rise of public [cloud computing](#) and its adoption by enterprises of all sizes is presenting challenges to professionals who are charged with the security of the organization's data. One major issue is that individual departments and even employees can purchase public cloud services—often by using a corporate credit card—without the knowledge or oversight of the IT department. Such purchases can lead to significant governance challenges, introduce unknown risks, and even prevent the organization from meeting its statutory and regulatory compliance obligations.

Public cloud computing is desirable for many reasons, including increased IT agility, reduced time necessary to roll out a new product or service, access to the latest technology not available inside the enterprise—and even a strategy to work around restrictions put in place by the IT departments, such as a limit to the size of email attachments or the types of files that can be sent or received through the email system. For these reasons, many IT departments are considering deploying private clouds, which departments can access and use instead of public clouds. Examples on record include State Street Bank (which expects to see significant savings as well as improve

operational efficiency and security of customer data), engineering and construction firm Bechtel Corporation, and chemical company Sinochem Group.

However, private clouds aren't inherently more secure than public clouds and can even be far *less* secure. In this article, I'll discuss some pitfalls and make recommendations for securing private clouds.

Overview of Private Clouds

One difficulty that security professionals encounter is a variety of perceptions amongst IT staff, senior management, and end users about what a private cloud really is. For example, many believe that private clouds are exclusively on-premises, residing in a data center that the organization controls, and therefore are more secure than a public cloud, which is hosted by a public cloud provider. Another common misconception is that a private cloud always uses virtualization to create a pool of virtual machines (VMs), which can be allocated to departments and users as needed.

Although it's fair to say that most private clouds that are deployed or under consideration today are on-premises, provide Infrastructure as a Service (IaaS), and use virtualization technology to create a pool of resources that can be allocated as required, the reality is that a private cloud is simply a cloud that is dedicated to an organization for its sole use. The private cloud can be on- or off-premises, and indeed several extremely reputable private cloud providers use their own data centers to host private clouds for their customers. Some of these private cloud providers (e.g., Microsoft) also offer public cloud services. Nor are private clouds limited to IaaS offerings. Many can and do offer Platform as a Service (PaaS) or Software as a Service (SaaS) as well. Virtualization, especially in SaaS clouds, is not a prerequisite. Organizations might find it useful to consult the National Institute of Standards and Technology's (NIST's) Special Publication 800-145 for a better understanding of cloud computing, including private clouds. (See the sidebar "Cloud Definitions" for more information.)

Cloud Definitions

The National Institute of Standards and Technology's (NIST's) Special Publication 800-145, "[The NIST Definition of Cloud Computing](#)," defines the cloud as having five essential characteristics:

- On-demand self-service
- Broad network access
- Resource pooling
- Rapid elasticity
- Measured service

The publication also specifies three service models:

- Software as a Service (SaaS)
- Platform as a Service (PaaS)
- Infrastructure as a Service (IaaS)

And the document defines four deployment models:

- Private cloud
- Community cloud
- Public cloud
- Hybrid cloud

The NIST standard is freely available for use and remains the best definition of cloud computing today, in part due to its simplicity and brevity. NIST has published other cloud standards that are primarily aimed at departments and agencies of the US Government, as well as cloud architecture document Special Publication 500-292, "[NIST Cloud Computing Reference Architecture](#)." Larger enterprises might find this document useful when planning adoption of the cloud. ■

InstantDoc ID 144176

Threats in Private Clouds

The [Cloud Security Alliance \(CSA\)](#) published its [Top Threats to Cloud Computing](#) research report in March 2010. Although the document is several years old and is currently being revised, it remains extremely relevant today. Of the seven threats identified in the report, each applies to private clouds, whether hosted on- or off-premises. All seven apply to IaaS, six apply to PaaS, and five apply to SaaS clouds. The threats, in no particular order, are as follows:

- Abuse and nefarious use of cloud computing
- Insecure APIs

- Malicious insiders
- Shared-technology vulnerabilities
- Data loss or leakage
- Account, service, and traffic hijacking
- Unknown risk profile

The CSA has published guidance and tools that cloud consumers and providers can use to jointly tackle these threats. These tools are freely available on the [CSA's website](#).

The reported threats aren't considered the only threats to cloud computing, though. Experience with private clouds has highlighted specific areas of concern for enterprises. As I mentioned earlier, most private clouds provide IaaS and use virtualization to make VMs available to departments and users within the organization, on an as-needed basis. Both Microsoft and VMware provide the necessary technology to build private clouds from the ground up. Excellent open-source tools, such as OpenStack, work with a range of hypervisors such as Citrix Xen and Linux Kernel-based Virtual Machine (KVM), as well as Microsoft and VMware hypervisor technologies. For that reason, I will focus on IaaS risks for the remainder of this article.

Risk #1: Abandoned VMs

Private clouds have led to an explosion in the number of VMs in existence. Private clouds are used to develop and test line of business (LOB), or Tier-1, applications and customer-facing web-based applications, as well as to host production environments. Organizations often develop entire libraries of VMs that can be deployed at a moment's notice, to handle additional workloads or to accommodate specialized testing. Creating new VMs can, in some cases, be as easy as copying the configuration files that define a VM and the file or files that comprise its virtual hard disk (VHD). This explosion has led to a problem whereby VMs are created and used but rarely deleted. When a VM is no longer necessary, it is often simply turned off and left intact, just in

case it ever needs to be used again. This approach might mean placing it back in a library. This phenomenon is facilitated by the relatively low cost of storage used to house the VM libraries.

The risks to the enterprise from such practices are many. For example, VMs that are simply turned off typically are not turned back on to apply routine software updates. When a VM is used only during periods of peak demand, it might go weeks or even months between uses, by which time it might have several critical vulnerabilities for attackers to exploit as soon as it comes online.

Every organization should put in place a set of procedures to ensure that VMs have all software updates applied before being used in production. Unfortunately, many organizations do not create a production-like environment in which the VM can be brought online, updates can be applied, and then regression can be tested before the VM is moved into production. Either these organizations haven't planned for such an environment, or they don't have the time or resources to create one. When in production, each running VM should be treated just like a physical server and should have all appropriate software updates, maintenance releases, and so on applied, to ensure that no vulnerabilities can be exploited by attackers.

Risk #2: Exposing Sensitive Data

Another common problem with the explosion of VMs in private clouds is the amount of highly sensitive data that is stored on these VMs. This sensitive data can include personally identifiable information (PII) such as usernames, passwords, Social Security numbers, or credit card details. It can just as easily hold protected health information (PHI) or the organization's business secrets, including source code and product plans. When a VM is running in a production environment, access to a live VHD is tightly controlled by the hypervisor or other virtualization software or hardware. The VHD file (or files) is often locked for exclusive access by the hypervisor, to prevent corruption. The OS in the running VM implements logical access control,

preventing access to sensitive data just as an OS would do on a physical server when users attempt to access the server over the network or to interactively use a remote command-line or desktop interface.

Plenty of tools can be used to mount a VHD as a drive letter or mount point when not in use. These tools can allow an attacker unrestricted access to the data on the VHD, including extremely sensitive data. However, malicious attackers aren't the only ones accessing offline VHDs, looking for sensitive information that organizations need to be worried about. Private clouds and virtualization make it easy for developers to get access to production servers that are, in actuality, virtualized servers, for the purposes of testing and debugging simply by copying the configuration and VHD files to their development environment. There are many restrictions in place about access to various types of data, and the organization might need to disclose, to regulators and even to customers, the inadvertent access to data by developers. This disclosure can bring the prospect of fines or other penalties, and the loss of customer goodwill and business. For these reasons, access to the configuration and VHD files that make up VMs should be restricted and monitored. Although it comes with a performance overhead, sensitive data should be encrypted when stored in a VM, and the encryption keys should be stored elsewhere, such as a network-based encryption device (e.g., a Hardware Security Module—HSM). When it is impossible to encrypt data stored within a VM (e.g., when running a commercial off-the-shelf—COTS—application that doesn't support encryption in a VM), it might be possible to encrypt the VMs' configuration and VHD files, depending on the virtualization technology used. Last, logical access controls such as discretionary ACLs (DACLs) should be used in the host environment and in the library, to restrict access to the VM files when not in use.

Risk #3: No Network Access Controls

A third problem commonly encountered in private clouds, especially in private clouds where VMs can be moved from host to host to allow

for maintenance of host machines or to ensure uptime in the face of host failure, is that the networks in which the machines run are relatively flat, perhaps using RFC 1918 Class A addresses in the 10.x.x.x or Class B addresses in the 172.16.x.x – 172.31.x.x ranges. Flat networks make for more-simple-to-build and easier-to-manage physical networks connecting the VMs. These networks also help to ensure connectivity from one VM to another, to a physical machine, or to the Internet, regardless of the physical host on which the VM is running.

Unfortunately, flat networks rarely have much, if any, security in place. There are fewer routers or firewalls on which to place ACLs, and there's no single point at which Intrusion Detection Systems (IDSs) or Intrusion Prevention Systems (IPSs) can be deployed to monitor all network traffic. The lack of ACLs and IDS or IPS means that as soon as an attacker or piece of malware is in the network, it can move around with relative ease and with little chance of detection.

Although network-level controls are preferred and even a best practice as part of a defense-in-depth strategy, the lack of security in the network can be compensated somewhat through the use of host-level firewalls within the VMs. Unfortunately, configuring these firewalls can be a management nightmare. There is a lack of tools for the task, and many VMs reside offline in a library, and so cannot be easily managed. Some virtualization software does allow you to configure network access to VMs and comes with management tools to support configuration and maintenance. And several industry alliances promote best practices in network security through reference architectures and tools for networking components and virtualization software from specific vendors. However, these are proprietary and won't work in moderate to heavily heterogeneous host environments.

Risk #4: Lack of Anti-Malware

Another security problem that is routinely encountered in private clouds is a lack of anti-malware defenses. There is much confusion and misunderstanding about the role of anti-malware software in

IaaS clouds—a holdover from when virtualization first became mainstream. For example, many organizations won't employ it at all on the host or in guest VMs, believing that it's either not required, will cause performance issues, or might even cause a VM to fail.

The truth is that anti-malware defenses *are* required and need to be carefully deployed to ensure smooth operation and protection of host machines and guest VMs. Careless deployment of anti-malware software to a host machine can cause significant performance effects and can cause a VM to be flagged as malicious (either because it's infected or because the malware signatures of the anti-malware software in the guest VM are detected and flagged as malware).

Anti-malware software on the host machine should be configured to exclude the scanning of files that comprise VHDs and the processes that are associated with the virtualization software. Anti-malware software in VMs needs to be configured so that it doesn't cause performance problems on the host machine, which could affect performance across other guest VMs running on that host. This requirement might mean configuring the anti-malware software in VMs to scan for malware in files during off-peak hours, and limiting the amount of CPU and memory resources that the anti-malware software uses.

Virtualization best practices, including dedicating a CPU core to each VM and placing each VM's configuration and VHD files on separate physical SCSI hard disk spindles, will largely mitigate the possibility that anti-malware software in a VM will have a performance impact on the host machine and other guest VMs. As with applying software updates, which I described earlier, the organization needs to ensure that offline VMs in a library are routinely brought online to obtain the latest anti-malware software and signatures. Otherwise, when brought online, the VMs might fail to detect newer malware.

Risk #5: Decentralization

The final security concern that I want to highlight is a lack of understanding about how departments and individuals use private cloud

resources. When an organization deploys a private cloud, there is often a tendency to decentralize the IT department, assigning IT staff to business units and departments that might also hire their own IT staff, architects, and application developers. The role of the IT department descends into one whose primary purpose is simply to keep the private cloud infrastructure up and running and to allow the business units and departments to manage their own VMs. In such environments, business units can begin to make compliance- and risk-based decisions about what to put in the private cloud and which controls and protections to put into place. Unfortunately, these individual business units and departments, as well as whatever is left of the central IT department, are unaware of other business units' and departments' use of the shared private cloud. As a result, they might introduce an unacceptable level of risk to the entire organization.

For example, suppose that one business unit deploys a customer-facing application that collects and processes credit card information, taking care to ensure that the appropriate controls are in place. That effort can be imperiled by another business unit that deploys a customer-marketing website for a seasonal campaign but does *not* deploy suitable security controls for an environment that processes credit-card data. Or suppose that a business unit or department deploys an application that requires high security because it processes customer PII. However, VMs are stored in a library, with little to no protection, when not in use.

To ensure that no business unit or department introduces unnecessary risk, a centralized function must understand which groups are deploying what into the private cloud, ensure that the private cloud infrastructure can support all security requirements, and ensure that no group is deploying insecure applications. To accommodate environments in which differing business units have differing security requirements, an organization can choose to either partition the private cloud into high-, medium-, and low-security segments, or can consider deploying more than one private cloud.

Balance the Benefits and Risks

Private clouds can bring many benefits, including reduced costs, increased business agility, the ability to empower individual business units and departments, and a closer alignment of IT and business needs. But private clouds can also introduce new risks. Organizations that are considering private clouds—as well as those that have already deployed them—should consult guidance such as that made available by the CSA and the software manufacturers whose products and tools the organizations are using to build out their private cloud environments. Other resources that might be of use are those made available by public cloud providers such as Google, Amazon, Microsoft, HP, Verizon, and Rackspace. These resources are intended primarily to address the security and privacy concerns of potential customers but also contain a wealth of detail about how these vendors built out their own cloud infrastructures and the types of controls that they put into place to secure customers' data. ■

InstantDoc ID 144175

The Essential Guide to Achieving High Availability for SharePoint Data

The growth of SharePoint around the world has exposed one of the major infrastructure challenges of the application: that there is no inherent two-way data replication engine built into the tool. Because of this, organizations face a choice between centralizing their SharePoint data and having users access it across high-latency or low-bandwidth links, or dealing with a decentralized SharePoint content structure, which often results in duplication of data, outdated content, slow end user performance, operations outages, regulatory risks, and discontented users.

This essential guide focuses on understanding the various approaches organizations have taken in dealing with the explosion of SharePoint content across their organizations, and how to best distribute their content so that everyone can

work with the most up-to-date content in the most reliable, efficient manner. We'll focus on some of the built-in options such as new ones provided within SQL Server 2012, and also compare those options to third-party replication approaches, including information on how these approaches can be used to achieve high availability and business continuity objectives.

Understanding SharePoint High Availability Challenges with SharePoint

SharePoint products and technologies have fast become a preferred method of providing for advanced document management for many organizations. Critical records and documents are stored within it, and are managed and made easily searchable within the application. As SharePoint envi-

WindowsITPro

Metalogix

ronments become more widely used, their importance increases exponentially as well. No longer is SharePoint simply a “dumping ground” for old documents. In many organizations SharePoint has become a mission-critical application for full enterprise collaboration.

As such, mission-critical applications typically need to be built to be both fault tolerant and highly available, and SharePoint is no exception. Building high availability into SharePoint has multiple challenges, however, because it is a complex application with various operation tiers that need to be factored into the discussion.

High availability does not just involve building redundancy into the SharePoint hardware. It also involves architecting a SharePoint environment that allows for relevant data to be close to the users that need to access it, and providing for multiple copies of this data to provide for failover in the event of an outage. The native architecture of SharePoint and the fact that a single SharePoint farm can’t easily span across multiple sites can make this a challenge, however.

The first challenge with high availability design for SharePoint has to do with the fact that SharePoint itself operates at three very distinct layers, each with their own operational requirements, as illustrated in Figure 1.

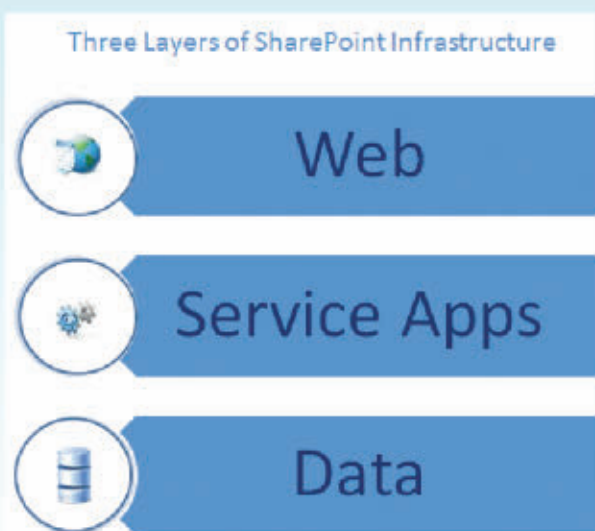


Figure 1 – Three Layers of SharePoint Infrastructure

Exploring Native High Availability Options with SharePoint

The Web Tier handles the traffic from the users, and it runs on Internet Information Services (IIS). It translates all of the user-based traffic directly into the platform, allowing users to browse pages, upload and download documents, and view reports. High availability at the Web Tier typically involves the use of multiple servers in a SharePoint farm all running as Web Front-end Servers. A load balancer distributes user requests to these servers to optimize response times and offer redundancy in the event of a server failure. While in many cases this may be hardware-based load balancers, in some cases, this may also in-

volve the use of a software load balancer such as Windows Network Load Balancing, included in the Windows Server operating system. Windows Network Load Balancing can be a challenge to implement, however, and it is often best to consider the use of a good hardware-based load balancer, many of which are getting considerably cheaper over the years.

The middle tier of SharePoint is known as the Service Application Tier, and it stores shared services such as Search, the Managed Metadata Service, Excel Services, the User Profile Sync Service, and much more. This tier is the most diverse, because each Service Application has its own operational requirements and the majority of them have their own unique high availability concerns. For the most part, high availability at this tier can be achieved simply by running a service application on more than one server in a farm, though the major exceptions to this include the User Profile Sync Service Application and the Search Administration component. But, in general, if you set up at least two servers in a farm, and turn on each service application on both servers, SharePoint will automatically use both servers. In the event of an outage, the second server will then continue to operate, preserving the service application functionality for that farm.

The Search service application is one that requires some additional thought into high availability architecture, however. Simply turning on the service on two servers will not provide for automatic high availability. Instead, you need to make use of the ability to create multiple Index Partitions, query components, and crawlers and split them across the servers running the service applications, similar to what Figure 2 shows.

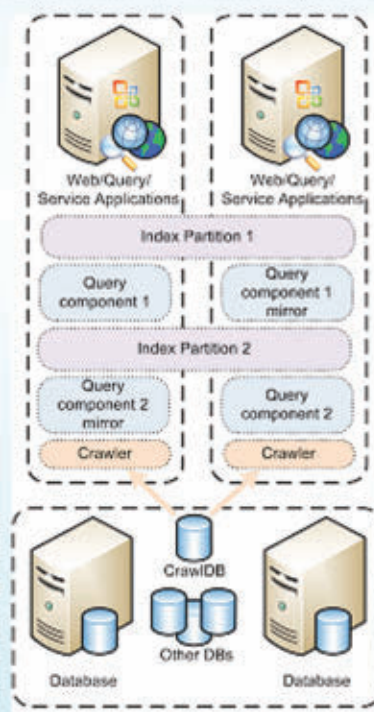


Figure 2 – Search Service Application High Availability

The final tier of SharePoint is the Data Tier, which runs on Microsoft SQL Server. SharePoint itself is highly dependent on this tier, as

all information stored in a SharePoint site is stored in SQL Server databases. Indeed, not only the content itself but also most of the information stored by the various Service Applications is stored on the SQL Server. High Availability at this tier has traditionally taken place via 'traditional' shared storage clustering models that involves two or more SQL Servers that share a common set of databases, through the use of SQL Server 2012.

Using new SQL Server 2012 High Availability Options such as AlwaysOn Availability Groups for SharePoint

The biggest recent change to high availability design for SharePoint farms is related to the recent release of Microsoft SQL Server 2012. The options provided with this version of SQL Server (Table 1) are significantly more sophisticated than some of the prior options, and provide for some fascinating new design options for SharePoint environments.

As indicated in Table 1, vastly different levels of Service Level Agreements (SLAs) are provided by the various Data Tier options. Traditional Backup and Restore options, for example, may end up with data loss that totals hours or days, and may take weeks to bring back online, while some of the newer options provide for zero data loss and failover within seconds. Subsequently,

it is critical to carefully consider what the preferred Data Tier availability options for an organization will be.

The most significant new feature added in Microsoft SQL Server that changes the design paradigm for SharePoint administrators is a feature known as AlwaysOn Availability Groups (AOAGs). AOAG technology is a combination of SQL database mirroring technologies, where exact copies of databases can be made on other servers, together with clustering technologies that allow for automated failover. By combining these technologies, AOAGs allow SharePoint design architects to create multiple redundant copies of SharePoint databases, allowing for up to five total copies of the content.

Key features for AOAGs for a SharePoint environment include the ability

- to create synchronous replicas of SharePoint databases on up to three servers, with automatic failover between two of the replicas,
- to create up to four asynchronous replicas of SharePoint content databases, providing for the ability to direct content to remote sites,
- to set some of the replicas as 'read-only' replicas, creating scenarios where that data is connected to a 'Read Only SharePoint farm' in a remote location,
- to provide for automated failover within five to seven seconds, of all SharePoint content in a farm.

High Availability Options for SQL Server	Potential Data Loss (RPO)	Potential Recovery Time (RTO)	Auto-Failover	Additional Readable Copies	Limitations for SharePoint HA
AlwaysOn Availability Groups-Synchronous (Dual-phase commit, no data loss, can't operate across WAN)	None	5-7 Seconds	Yes	0-2	Cannot work across low bandwidth/high latency networks
AlwaysOn Availability Groups- Asynchronous (Latency tolerant, cross WAN option, potential for data loss)	Seconds	Minutes	No	0-4	Cannot be used for SharePoint Service Application and/or Config databases
AlwaysOn Failover Cluster Instance (FCI) - Traditional shared storage clustering	N/A	30 Seconds to several minutes (depending on disk failover)	Yes	N/A	Cannot work across low bandwidth/high latency networks. Only one copy of database files.
Database Mirroring - High-safety (Synchronous)	None	5-10 Seconds	Yes	N/A	Cannot work across low bandwidth/high latency networks
Database Mirroring - High-performance (Asynchronous)	Seconds	Manually initiated, can be a few minutes if automated	No	N/A	Cannot be used for SharePoint Service Application and/or Config databases
SQL Log Shipping	Minutes	Manually initiated, can be a few minutes if automated, by typically hours	No	Not during a restore	Cannot be used for SharePoint Service Application and/or Config databases
Traditional Backup and restore	Hours to Days	Typically multiple hours, days or weeks	No	Not during a restore	Not a high Availability technology

Table 1 – Comparison of High Availability Options for SQL Server

This creates new design options for SharePoint farms, which are illustrated in Figure 3.

Figure 3 illustrates a sample design for AOAGs that takes full advantage of all of the features of the technology. In this ex-

ample, a primary datacenter is configured for automated high availability of the SharePoint Data Tier, with synchronous copies of all databases configured to automatically failover within five to seven seconds between the primary SQL server and

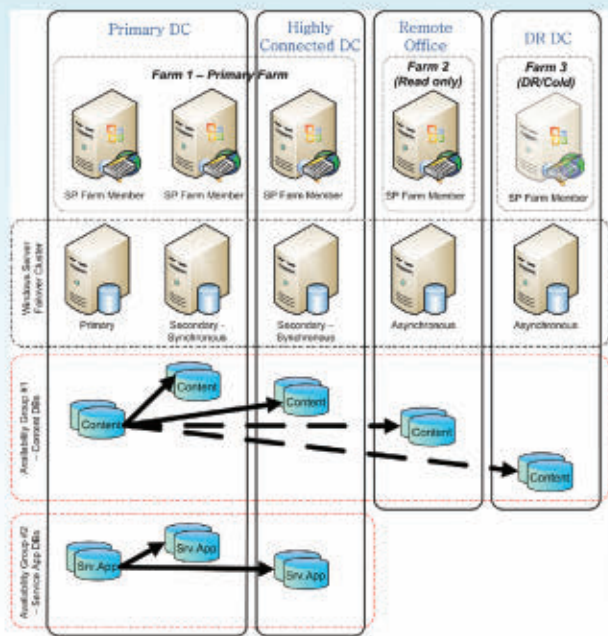


Figure 3 – AOAG Design Options for SharePoint

the secondary replica. In addition, a third synchronous copy of the content is set for manual failover to a highly connected datacenter located close by. For synchronous AOAGs to work properly, the connectivity must be very high (1 GB or greater) and the latency must be very low (10ms or less.)

In this example, two additional asynchronous replicas of content database are created in remote datacenters. The reason that only content databases have these fourth and fifth replicas created has to do with the fact that only the Secure Store database and the content databases of SharePoint are supported for asynchronous replicas;

all other SharePoint database types support only Synchronous replicas. In this model, the first asynchronous replica is created in a remote site that is configured with a Read-Only SharePoint farm that is configured to allow users in that location to have a local readable copy of their SharePoint content. The final datacenter has an asynchronous copy put in place to provide for a location to failover the SharePoint content to in the event of a disaster.

This model illustrates an example of how you can use the SQL AOAGs to improve some of the design options, and provide for levels of local high availability that were previously unattainable. However, with every new solution comes a set of limitations, which are important to note for AOAGs. These include the following:

- While the read-only replicas allow for a remote site to have faster access to SharePoint content, a separate URL must be given to the users to access the read-only copy vs. the full read-write copy in the home office. This can lead to confusion, because user's must understand to go to <http://readonly-sharepoint.companyabc.com> when they want fast read access to SharePoint but go to <http://readwritesharepoint.companyabc.com> when they want to make changes to content.

- Using AOAGs with SharePoint require the most expensive Enterprise editions of both SQL Server and Windows Server 2008 R2 or 2012.
- Service Application functionality can only be replicated synchronously; the asynchronous replicas in remote locations cannot have the critical service application functionality replicated to them.

Examining Third-Party Replication Models

Because of the limitations inherent in SharePoint high availability options, many clients over the years have opted for approaches that use application-layer replication technologies to be able to provide for high availability for SharePoint content. These technologies typically operate at the SharePoint API level, looking for new documents and content within a SharePoint site, and then sending that content to another active farm in a different location. Every time a document is added or modified in SharePoint, a copy of that document is then replicated to the multiple farms within the organization.

In many cases, these multiple farms are configured to take advantage of global load balancing modules that allow for a single URL to be used for all of the farms. That way, if a user clicks on a link to a SharePoint site, the load balancer determines what the

closest replicated farm is and sends the user to that farm. Since all content is kept in sync via the third-party replication tools, this allows for a seamless experience for the end users, and keeps content closer to the users, while allowing for multiple live copies of the content to exist in multiple locations. For a high availability perspective, if a local farm is down due to an outage or for maintenance, the load balancers can also send the users to a different farm in a remote location, preserving availability. These options have been successfully implemented in various forms throughout the years, and are especially interesting from an architecture perspective as they provide for concepts unavailable with the out-of-the-box tools available with SharePoint or SQL Server.

Think Through the Process

Designing for high availability in a SharePoint farm can be a complex process. SharePoint architects need to take into account availability at the Web Tier, the Service Application Tier, and the all-important Data Tier. New options in SQL Server 2012, including AlwaysOn Availability Groups, can improve an architect's design options, while other organizations may choose to look at the flexibility and robust high availability models that are provided by third-party replication models instead. In any

case, because of the importance of SharePoint in most organizations, proper thought and care should be put into the design and implementation of high availability for SharePoint farms.

ABOUT THE AUTHOR

Michael Noel, is an internationally recognized technology expert, best-selling author, and well known public speaker on a broad range of IT topics. He has authored several industry books that have been translated into over a dozen languages, with sales exceeding 500,000 copies worldwide. Significant titles include *SharePoint 2013 Unleashed*, *Windows Server 2012 Unleashed*, and *Exchange Server 2013 Unleashed*. Currently a partner at Convergent Computing in the San Francisco Bay Area, Michael's writings and extensive public speaking experience across over 100 countries and all seven continents leverage his real-world expertise helping organizations realize business value from Information Technology infrastructure.

Best Practices for Group Policy Design

Use these guidelines for optimal performance and security



**Darren
Mar-Elia**

is a Microsoft MVP for Group Policy, a contributing editor for *Windows IT Pro*, and CTO and founder of [SDM Software](#). He maintains a Group Policy resource website and has authored many books on Group Policy and Windows topics.

Email



Website



Last June at TechEd North America, I presented a talk on best practices for Group Policy design with performance and security in mind. This article is based on that talk. In it, I hope to provide you with some guidance and best practices on what makes good Group Policy design.

What Is a Good Group Policy Design?

Before you can aim for a best practice, you need to have some idea of your target. In the context of Group Policy, what are good design criteria? I typically shoot for achieving a balance of the following elements:

- Minimal impact to the end user
- Balance of security and lockdown goals
- Minimal management overhead and complexity

The challenge in all of this is that sometimes these three goals work at cross purposes to one another. For example, having a minimal impact on the end user while also meeting your organization's security goals might be difficult.

When I first started using Group Policy (and even its predecessor, System Policy in Windows NT 4.0), I had a tendency to take the newfound power I'd gained with these technologies and tweak as many switches as possible. Over the years, I've learned that more isn't necessarily better when it comes to Group Policy. That's why it's also

important to know your business needs before you embark on tightening the proverbial policy screws.

Another key question that I'll dive in to is how many Group Policy Objects (GPOs) are too many. And once again, I take the "Goldilocks" approach to this question—just enough GPOs and just enough complexity to accommodate the needs of the business yields a design that's just right.

OK, enough about the "soft" side of Group Policy design. Let's talk nuts and bolts!

Monolithic vs. Functional GPOs

One decision you'll need to make as you deploy GPOs in your environment is how you'll group settings within those GPOs. I use the terms *monolithic* and *functional* to describe the two possible approaches:

- Monolithic GPOs—Contain a variety of settings from multiple policy areas (e.g., Administrative Templates, security, Group Policy preferences)
- Functional GPOs—Contain one or more settings from a single policy area and often target a single function (e.g., Internet Explorer—IE—lockdown)

Most environments have a mix of both types of GPOs—driven by factors such as the need for delegating certain GPOs to a particular business unit administrator, the desire to manage complexity, and the need to enforce security mandates. A perfect example of a functional GPO is one that contains domainwide security policies. In that situation, you might have one functional GPO that sets those domainwide policies and can be edited by only one person or group (e.g., the security operations group).

The goal of the functional GPO is to isolate one or more settings from a single policy area so that they can be handled as a unit and easily delegated to a particular user or group if need be. However, you can go overboard with functional GPOs. Having 100 GPOs, each with

**A good
Group Policy
design should be
"seen but not
heard."**

a single setting, is a tangible example. There are performance penalties to be paid if you have many, many GPOs, though those penalties might not be as bad as you'd expect.

The goal of monolithic GPOs is to bundle together a complete configuration scenario within one GPO. Be it the Marketing Desktop Lockdown GPO or the Mobile User Configuration GPO, the goal is to keep all settings that relate to the scenario in one manageable and delegated GPO. Monolithic GPOs are ideal for the organizational unit (OU) administrator who needs to have control over Group Policy settings for users but shouldn't be able to create GPOs willy-nilly. Monolithic GPOs also ease troubleshooting: Because one GPO is responsible for the lion's share of policy, you have only one place to look when something goes awry.

The bottom line is that monolithic or functional GPOs both make sense in certain situations, but one might be better than the other from a Group Policy processing-performance perspective. More on that later.

Linking vs. Filtering

The question of where to link your GPOs and when to use filtering on them is another decision that has both complexity and performance ramifications for your GPO design. You're faced with two choices:

- Link GPOs as close to the intended target as possible.
- Link GPOs higher in the Active Directory (AD) hierarchy, and then rely on filtering (e.g., security groups, Windows Management Instrumentation—WMI—filters, Group Policy preferences item-level targeting) to get the setting where it needs to go.

You'll find reasons to support both choices. Much of the decision will depend on your AD design (which I'll talk about in the next section). My guiding principle is that you should always seek to link a GPO as close to the intended targets as possible and rely on filtering on an exception basis only. This is easier said than done when you have

a flat AD hierarchy (i.e., all users in one OU). But if you can adhere to this principle, you not only reduce complexity in your environment, you also take away the performance impact of needing clients to evaluate a variety of filters on a given GPO to determine whether it applies. And if you throw in the use of Group Policy preferences item-level targeting, which lets you create filters on individual settings, well . . . it's easy to see how things can get out of hand. Pretty soon, your clients are evaluating tens of filters just to determine whether a given GPO applies to them, let alone processing GPO settings. The performance impact can be even more profound if you're using item-level targets that require network communication to resolve. Examples of these targeting items include Security Group, LDAP Query, Domain, Site, and Organizational Unit, all of which require LDAP calls to AD.

Again, using linking as the primary mechanism for targeting GPOs is a good practice. Filtering is definitely a valuable feature of Group Policy and Group Policy preferences, but be aware of the possible performance and complexity penalties that you can pay if you rely too heavily on those targeting tools.

Balancing Active Directory and Group Policy Design

When it comes to how your AD design interacts with your Group Policy design, finding common ground is often a struggle. AD designs tend to be driven by criteria such as application requirements, delegation, and administration. Group Policy is driven by targeting convenience, platform type (i.e., server vs. desktop), or security goals. That said, it's important to consider Group Policy needs when you're designing (or redesigning) AD. Keep these design goals in mind during your AD design discussions:

- Try to deploy 80 percent of your GPOs without requiring filtering.
 - Find an OU design that lets you link close to the target for 80 percent of your scenarios.
 - The other 20 percent should require compromises, not AD redesigns.

- Avoid designs that force you to link and enforce at the domain level. Such designs make downstream changes more difficult. Reserve linking at the domain level for truly global settings, which should be few and far between.
- Avoid overly flat OU structures (i.e., all users in one OU) if you plan to use per-user policy in any significant way. Such structures require massive reliance on security filtering, which adds tremendous complexity and risk (e.g., large numbers of objects might be mistargeted if the wrong filter is applied).
- Avoid designs that require loopback for all computers. Loopback should be reserved for scenarios such as kiosk systems, Remote Desktop Services servers, and so on.

Now that we've talked about high-level principles for designing and deploying GPOs, let's dig into some of the technical aspects around how GPOs are processed and how different deployment decisions can affect Group Policy processing performance.

Understanding Group Policy Processing

A big part of the decision process for designing and deploying GPOs is understanding how Group Policy processing occurs under the best circumstances (and what happens when circumstances aren't optimal). To that end, let's talk about the different ways that GPOs are processed. Some of this might be review for you, but it's important foundational knowledge for understanding more complex scenarios.

Background vs. foreground processing. As you know, there are two types of Group Policy processing events: foreground and background. For computers, foreground processing happens on startup. For users, processing happens at logon. Background processing, as the name implies, occurs periodically, based on the client's role. Domain controllers (DCs) perform a background refresh every 5 minutes, whereas client OS versions and regular Windows servers perform a refresh every 90 minutes plus a 30-minute, randomized interval. In addition,

Windows Vista and later clients perform a background refresh based on network state. Specifically, if a Windows client (e.g., a roaming laptop) is out of contact with a DC when a background refresh is due, that client immediately performs a background refresh as soon as the DC becomes available. This refresh is often referred to as a network location awareness (NLA) refresh.

Synchronous vs. asynchronous processing. Another important aspect of Group Policy processing that has significant performance ramifications is the distinction between synchronous and asynchronous Group Policy processing. To understand synchronous processing, let's look at a typical example of Group Policy processing from bootup to user logon.

When a Windows computer starts, there's a point at which the client connects to the network. At that point, computer-based Group Policy processing kicks off. If this processing is configured to run synchronously, the user doesn't see the logon dialog box (aka the Graphical Identification and Authentication—GINA) until the processing is completed. After the user logs on to the system, user-based Group Policy processing begins; the user doesn't see the desktop until that processing finishes. Thus, synchronous processing elongates the time it takes for a user to boot up the system, log on, and get productive.

But starting with Windows XP, Microsoft set the default for foreground Group Policy processing as asynchronous. This type of processing is also called Fast Logon Optimization and remains the default foreground processing method through [Windows 8](#). Asynchronous processing basically tells Windows to continue doing what it was doing, even if Group Policy processing is still running. So, when a computer boots up, it doesn't wait for computer-based Group Policy processing to finish before presenting the user with the logon dialog box. Likewise, when the user logs on, there's no waiting on user Group Policy processing before presenting the user with the desktop.

Most folks read this and think, "Why would I ever want to run Group Policy processing synchronously?" The answer, as many of you

have likely discovered, is that some Group Policy client-side extensions (i.e., Software Installation, Folder Redirection, Disk Quota, and Group Policy Preferences Drive Mappings) work only when run synchronously. So, some folks essentially disable asynchronous processing to ensure that these policy areas do what they're supposed to do. These people enable the somewhat-mislabeled Computer Configuration\Policies\Administrative Templates\System\Logon*Always Wait for the network at computer startup and user logon* policy setting to force synchronous foreground processing.

The truth is, if you're willing to wait for a few computer restarts or user logons for these policy areas to take effect, you probably don't need to kill all the benefits of asynchronous processing by enabling this policy setting. The four policy areas that require synchronous foreground processing will signal to Windows to run synchronously the next time foreground system processing occurs, to ensure that they can process their settings. And I should mention: Background processing is, by definition, always asynchronous.

The Role of Change in Group Policy Processing

One big optimization that Microsoft has included in Group Policy from the very beginning is that, regardless of whether an event is processed in the foreground or the background, no processing occurs if nothing has changed within the GPOs that apply to a given computer or user. In such cases, Group Policy processing goes through the motions of reading all the GPOs that apply. But if, when comparing an existing GPO in AD with its record in the client's registry of what was done last time, the Group Policy engine notices that no changes have occurred, then each client-side extension that implements that policy area simply "walks away." Exceptions to this behavior can occur, such as when someone issues a `Gpupdate /force` command from a client. This command essentially says, "I don't care whether anything has changed—reprocess all policy anyway."

What constitutes a change that the Group Policy client cares about? Here are some types of changes that trigger a full reprocessing of policy:

- Someone makes a change that increments a GPO's version number. A difference in version numbers between the current, live GPO and the version number that the client last processed is considered a change.
- The list of GPOs that apply to a computer or user has changed. Causes for such a change include changes to security group filters on a GPO, changes to WMI filters that are linked to a GPO, or computer or user security group membership changes that cause a GPO to fall in or out of scope.

Note that even when changes trigger a full refresh of policy, not all client-side extensions completely refresh their settings. Let's say that you've deployed Microsoft Office via Group Policy Software Installation. Even if changes cause the Software Installation client-side extension to reprocess the GPO that delivered Office, that client-side extension isn't going to uninstall and reinstall Office. The extension simply reads the GPO's settings and makes sure that nothing major has changed (e.g., you've actually removed Office from the GPO).

The point is that if your environment is relatively static, the computer startup and user logon processes shouldn't be dominated by Group Policy processing. Processing should occur in milliseconds in most environments or in a few seconds in the largest environments.

Group Policy Processing Performance

The topic of Group Policy processing performance is always a touchy one. No one wants his or her GPO deployment to land on the CIO's top 10 list of reasons why users are unhappy with desktop performance. As a result, it's important to think about your Group Policy design in the context of having little impact on the user's desktop experience. To that end, let's talk about Group Policy behaviors and design decisions


that can minimize this impact. First, understand where time is typically spent during Group Policy processing, which is composed of two distinct phases: core processing and client-side extension processing.

During the core processing phase, the user or computer determines which GPOs apply and which client-side extensions must perform work. This phase is also when the Group Policy client determines whether something has changed and whether it needs to take action.

During the client-side extension processing phase, the hard work is performed. Each client-side extension that's registered on a system and that has work wakes up and process all the GPOs that were identified during the core phase. This second phase is when the actual settings are applied to the client system. In terms of proportion of overall time spent in Group Policy processing, the client-side extension processing phase easily consumes the greatest amount, as Figure 1 shows.

This pattern doesn't change much even when the client is processing many, many GPOs. It's just the nature of the beast. The time necessary to query AD for GPO information is generally much less than the time spent writing keys to the registry, installing software, mapping drives, and so on. Remember the question I brought up at the beginning of this article: "How many GPOs are too many?" To that point, the number of GPOs that you have is less important than what

Figure 1
Comparing Core and
Client-Side Extension
Processing Times



those GPOs are doing. If you have many GPOs that are each doing many things, then your GPO deployment will definitely affect the user's desktop experience. If you have many GPOs and half of them deploy only one registry policy setting, then the impact will be less. There's also the question of whether you want to maintain so many GPOs that are doing so little, but that's a separate consideration—one that I'd file under *managing complexity*.

There's also the question of how often your GPOs are changing. A ton of GPOs that rarely or never change aren't going to have much ongoing effect on your users, other than when a change does occur.

Assessing Client-Side Extension Performance

Now let's look at some low-hanging fruit around client-side extension performance. Examine the following client-side extensions, which are commonly examined around Group Policy processing performance:

- Client-side extensions such as Software Installation—which could be processing long-running software installations, or Folder Redirection, which copies user profile files around the network—take a long time to process the first time through. I'm not suggesting that you don't use these features. Rather, keep in mind that the first time they run, they'll usually have a heavy impact on users.
- The Security client-side extension—particularly using file or registry security against large trees of files or registry keys—can take a long time to process. This processing can have a significant impact, even during background processing, as the client-side extension churns through the resources that must be repermissioned. I typically avoid doing these kinds of large-scale permission-change tasks in Group Policy. A better method is to deliver the changes by using a one-time automation script or similar method.
- The Scripts client-side extension—more specifically, startup or logon scripts—are highly problematic from a performance perspective. Group Policy lets you have multiple scripts processing during a given startup or logon. That isn't necessarily a good

thing. Scripts have a tendency to linger in environments for years. Some run even though they long ago became irrelevant. Some network-intensive tasks access network resources that are no longer available. To top it off, most scripts don't have good logging, so troubleshooting delays caused by Group Policy-based scripts can be difficult. With the advent of Group Policy preferences, I typically recommend that, whenever possible, shops migrate common script tasks such as drive mappings, printer mappings, or simple registry tweaks to Group Policy preferences. It's a much more robust mechanism with a more complete troubleshooting infrastructure than you get with scripts.

Grouping Client-Side Extensions and Impact on Performance

Two distinct decisions that you can make about organizing your GPOs will have an important impact on performance.

The first decision relates to grouping frequently changing policy areas. Earlier, I alluded to possible performance impacts that relate to the decision to go with monolithic or functional GPOs. When you create monolithic GPOs that contain multiple policy areas, you might be inadvertently increasing Group Policy processing times. Why? The problem has to do with how the Group Policy engine detects a change to a GPO (which ultimately determines whether work must take place). That detection mechanism uses a simple version-number check on the GPO. So any change to a GPO requires that all client-side extensions that are implemented in that GPO must do work at the next processing cycle. Why? Because the Group Policy client has no way of knowing which policy area was changed in the GPO; it knows only that something changed.

To better illustrate this concept, let's use a concrete example. Suppose that a computer processes three GPOs: GPO A, GPO B, and GPO C. GPO A and GPO B implement registry and security policy settings. GPO C implements registry policy only. You decide to make

a change to security policy on GPO A. The next time Group Policy processing runs, it notices that the version number on GPO A has changed, but it doesn't know which policy area was changed. So processing must tell the registry and security client-side extensions that they both must process settings. In addition, even though GPO C has only registry policy implemented, the registry client-side extension must perform work, so it must process all GPOs within the computer object's GPO hierarchy. Processing only GPO A and GPO B would break that processing hierarchy. Figure 2 illustrates the process.

Suddenly, a simple change to one policy area in one GPO requires two client-side extensions to perform work across three GPOs. The moral of the story is that if you have frequently changing policy areas, grouping them together or putting them alone is better than mixing them with policy areas that don't change much. In the previous example, this approach would equate to moving the security policy areas out of GPO A and GPO B and putting them into their own GPO (or GPOs). Then, if you made a change to one of those areas, only the security client-side extension would need to work, and only against the GPOs that implemented the settings.

The second decision relates to our discussion about synchronous versus asynchronous processing. I mentioned four policy areas (i.e., Software Installation, Folder Redirection, Disk Quota, and Group Policy Preferences Drive Mappings) that require synchronous foreground

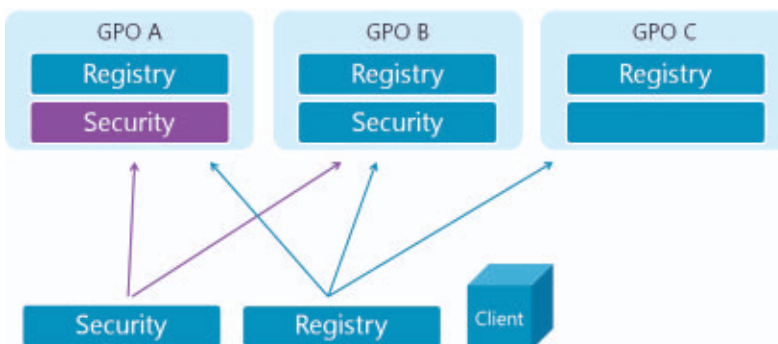


Figure 2
How Grouping of CSEs
Affects Performance

processing. Further, if any of these areas are implemented in a GPO and that GPO changes, then when any of these four client-side extensions process that changed GPO, they tell Windows to run the next foreground-processing cycle synchronously, even if the system is configured to run asynchronous foreground processing. And of course, if synchronous processing is configured, it elongates both machine-startup and user-logon times. Again, proper grouping of policy areas in GPOs comes into play. If you have a GPO that implements, say, Group Policy Preferences Drive Mappings and Registry policy, and you make a change to a registry policy setting in that GPO, then when the client processes the GPO, it doesn't know which policy area changed—only that a change happened. So both the registry and Group Policy Preferences Drive Mapping client-side extensions fire up. The Drive Mapping client-side extension tells Windows to run the next foreground cycle synchronously “just in case,” and suddenly an innocent little change to registry policy causes the next reboot or user logon to run slower! Just as in the previous decision around versioning and grouping of client-side extensions, when you're implementing one of these four synchronous policy areas, the best practice is to either put them in GPOs of their own or to combine them with each other, separate from policy areas that don't require synchronous processing.

Group Policy Performance and Loopback

It's worth saying a word about loopback and its potential impact on performance. As you probably know, loopback processing is typically used in kiosk, Remote Desktop Services, or Citrix XenApp environments. This type of processing comes in two modes: merge and replace. Merge mode has potential performance effects, depending on where your GPOs are linked in relation to computers that are enabled for loopback. That's because merge mode first processes user settings for the user object that's logging on to the loopback computer, then processes user settings that apply to the loopback computer object. This presents an interesting possibility: For example, the same GPO


containing logon scripts or other settings can actually process twice, if it's linked and filtered in such a way that it applies to both user and computer objects. The point here is that if you need or plan to use merge mode, use the Resultant Set of Policy (RSOP) modeling tools that come with Group Policy Management Console to determine what the effect will be on the user. That impact could literally double the time spent processing policy.

Seen But Not Heard

There are many things to think about when designing a well-performing Group Policy deployment. You'll need to consider everything from your AD design and where you link and filter your GPOs to how you group settings together. At the end of the day, a good Group Policy design should be "seen but not heard." Your users should not know that they're being managed by Group Policy. And you definitely don't want your CIO asking why Group Policy is causing so many problems. Plan, test, and deploy for optimal performance and security, using the guidelines that I've discussed here, and you—and your policies—will be in good shape. ■

InstantDoc ID 144059

Mobile App
Now Available!



Download your FREE mobile app.

iTunes | Android | Kindle

Business Benefits of Unified Communications


Take advantage of what UC has to offer

IT professionals have long been experts at deploying messaging systems. When email was a new technology, it brought about a huge cultural shift in businesses; now it's a well understood concept that both IT departments and end users have become accustomed to. Although a messaging deployment can touch a wide range of systems within a business, from the directory to the desktop, such a deployment can also be carried out without having a huge effect on day-to-day operations. Unified communications (UC) is a different type of solution, with multiple components that touch all aspects of a business.

Because a true UC platform includes many components, the number of touch points with IT systems, business systems, and users is greater than with other solutions. For example, in IT, UC affects the directory more broadly than an email project does, because of UC's reliance on a wider set of correct identity information (e.g., manager and telephone fields). A UC solution can also combine with video endpoints, telephone lines (and carriers), fax machines, and PBXs. From a business perspective, UC integrates with business applications and processes more widely than email does.

For end users, UC changes the way people work. Although email is now mature and widely adopted, UC brings new technology and a new mentality, providing the ability to truly change the culture and working practices of a business—which is a much more difficult task than the simple IT project that email has become.

Bringing UC into your business can present many challenges. In this article, I explain how a business might benefit from UC, and



Nathan Winters

is an Exchange technical specialist at Microsoft and leader of the UK Exchange Customer Council. He's lead author of *Mastering Lync Server 2010* (Sybex).



Email



Twitter



LinkedIn



Facebook



Website

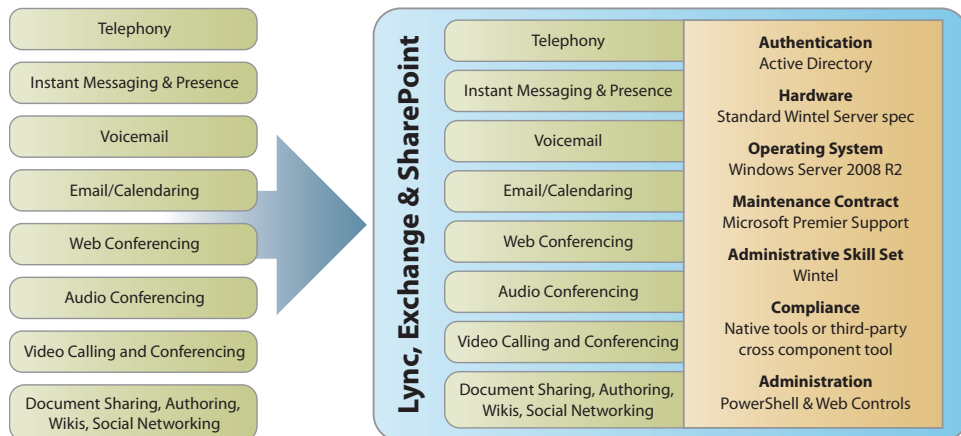
I discuss the various aspects of a UC project, including high-level technical components, the project team, different project phases, and common objections. I include resources that are available to help with your rollout. Finally, I cover the concept of cultural change, which in many cases is the most difficult hurdle to overcome but can provide the biggest benefit to your organization.

Benefits of Unified Communications

Most organizations will use several of the communication modalities listed on the left-hand side of Figure 1. Telephony, voicemail, and email are ubiquitous in some form or other; audio conferencing is also quite common. Other modalities such as video conferencing, web conferencing, and IM are becoming more widely used. The challenge for companies is that these technologies have often grown up in their business in very different ways.

For example, telephony and email have typically been implemented centrally, although often by different teams—but video conferencing might be present in only a small pocket or department. IM might be implemented through a public IM cloud such as Google Talk or Windows Live Messenger, over which a company has limited control. All these siloed communication technologies create overhead, such as the need to authenticate (thus remembering several usernames and passwords),

Figure 1
UC Technologies



look up names and addresses, and capture and archive business communication and documents for regulatory compliance purposes.

Security is also an important consideration; using a range of products for communication requires you to understand and secure a variety of systems. Implementing a unified platform lets you move toward a standard approach in which all information in transit—whether email or phone and video calls—is secured.

The right-hand side of Figure 1 illustrates the simplicity of bringing multiple technologies together in a single UC platform based on Microsoft Lync Server 2010, Exchange Server 2010, and SharePoint Server 2010. Implementing this type of solution gives you a single hardware platform, similar administrative interfaces and skill sets, one source of identity used for authentication and contact information (Active Directory—AD), and a more streamlined compliance platform—even though compliance systems aren’t entirely integrated for Lync, Exchange, and SharePoint out of the box. In addition, you empower users by providing a familiar and integrated user experience across all these technologies. The presence control and contact card lets users easily select the most suitable method for communicating with others, no matter which device they’re using or where they’re located.

In business terms, ROI is important. For UC projects, ROI falls into two areas: hard costs, which are easy to quantify, and productivity enhancements, which aren’t.

Hard cost reductions result from not having numerous maintenance contracts, not having multiple pieces of hardware and software, and not having several technical team members to manage a large number of unique systems. In addition, bringing conference bridges in-house that otherwise cost money for each minute used can greatly reduce hard costs. Finally, significant cost savings can be achieved through reduced travel expenses and the introduction of a more flexible work environment that makes better use of available real estate. (These last two changes imply the cultural shift that’s necessary for a company to truly benefit from UC technology.)

UC has multiple components that touch all aspects of a business.

UC has the power to change the way people work.

Productivity enhancements are more difficult to track and more difficult to attach a dollar value to—but they nevertheless result in noticeable improvements. For example, the ability to communicate efficiently in a variety of ways to suit the situation means that you'll typically be able to reach someone who can solve your problem. Whether your problem is a direct customer issue, an internal query, or even the simple need to get your expenses approved by a manager, taking care of tasks quickly can make a big difference in your productivity. A benefit of presence and integrated directory search is that you don't have to spend time chasing someone who isn't available; in addition, you can quickly find another person to resolve your problem by browsing the directory for someone in the same group (or that person's manager). These advantages, when combined with the empowerment that results from a flexible work environment (and the inherent trust in employees that's necessary in a flexible environment), contribute to increased employee satisfaction—which results in quicker issue resolution, leading to benefits such as a shortened sales cycle or better customer service. Another positive outcome is that you're more likely to retain the talented people in your organization.

Technical Components

Now that I've discussed the benefits of UC, let's take a look at what's necessary to have a fully integrated UC platform. The necessary elements can be broken down into different areas: Microsoft components, such as AD, Exchange, Lync, SharePoint, and Dynamics CRM; systems that interoperate with Microsoft's UC technology, such as video conferencing systems and PBXs; and internal factors, such as your network team and your security and compliance team.

Active Directory. AD is core to the success of a UC project. It provides a central point of identity, which enables authentication and helps locate users through directory search. AD provides the “single source of truth” that's so important in business. As such, AD should be populated with at least the following information about users:

- Job title
- Manager
- Office location
- Office and cell phone numbers in E.164 format (which is what Lync uses)

Exchange Server (including unified messaging—UM). Exchange provides a unified inbox, where all communication data is stored, including, email, voicemail, faxes, historic IMs, and even SMS messages. This approach gives you one place to search for communication information.

Lync. Lync is the glue that binds a UC system together. It provides the real-time communication components and enables significant integration not only into other Microsoft offerings such as SharePoint but also into third-party technologies such as room-based video conferencing systems and contact center systems.

SharePoint. SharePoint is where structured data is stored. In addition, SharePoint enables social communication in the form of My Sites and wikis that let people openly share information and skills, thus fostering community and tighter teamwork within an organization. SharePoint also enables the creation of custom workflows, which when tied into presence information can allow efficient routing of communication related to document submissions.

Dynamics CRM. Although often overlooked, Dynamics CRM can form an integral part of a UC system. Integration of presence information and communications modalities (both audio and IM) can enable efficient and knowledgeable servicing of customer queries.

Video conferencing systems. Systems from companies such as [Polycom](#), [Radvision](#), and [LifeSize](#) can all be integrated into Lync. Third-party video conferencing systems let Lync clients participate in meetings hosted on room-based conference bridges (and vice versa in some cases). When a technology such as Lync is fully deployed and integrated with other video platforms, desktop video capabilities

**Bringing UC into
your business can
present many
challenges.**

give users the ability to fully participate in and feel part of a conversation. In general, room-based systems have typically been reserved for senior executives. Integration with Lync means that users can take advantage of desktop video conferencing to participate in meetings, thus saving money and increasing efficiency. In addition, as video uptake becomes ubiquitous, the use of room-based conferencing systems will likewise increase.

PBXs. Organizations use a wide variety of PBX systems, often of varying ages and with widely different capabilities. Lync can integrate with these systems in a couple of different ways: as just another PBX communicating directly via Session Initiation Protocol (SIP) or through a gateway device, or through the use of remote call control or a similar technology. Remote call control allows Lync to be used for many capabilities and lets existing PBXs integrate with presence information and maintain a hold on telephony. Although this solution isn't necessarily ideal, it does merit investigation in organizations where there's an extensive or recent investment in PBX technology. At the very least, it allows Lync to become the centerpiece of the UC platform and to gradually take a greater role as existing assets expire.

Network. No matter which communications modality you intend to use, the traffic that's generated will flow across your network. Networks—and particularly the teams that run them—are critical to the success of UC adoption. It's important to involve the network team in planning a UC project from the early stages. Doing so can prevent arguments about bandwidth and can allay fears of “clogging up the network with so much rich media.” If you can bring the network team up to speed and verify the bandwidth required to implement a UC solution, you can have a discussion about Quality of Service (QoS) and Call Admission Control (CAC) in a rational setting rather than under stress. (For information about QoS and CAC in UC, see the sidebar “QoS and CAC in Unified Communications.”) It's likely that additional bandwidth will be necessary to realize the

full benefit of your UC solution, which might require creative budgeting. Therefore, it's helpful to have a senior project sponsor who understands the business benefits of UC and can evangelize UC on your behalf.

Security and compliance. As with the network team, security and compliance is another area in which you might encounter pushback. It's within a security professional's nature to say no, because it's his job to protect the organization from threats. Instead of springing changes on the security team, enlist their support early on. Again, use senior sponsorship to ensure that the security team understands the benefit that the UC project brings to the business. Although UC does lead to more open sharing of information, that sharing isn't uncontrolled or unmonitored. In fact, when they're implemented correctly, Lync and UC as a whole are extremely secure, with all server-to-server communication and server-to-client communication encrypted by default. In addition, all communication can be centrally monitored and archived for regulatory and compliance purposes—which is a vast improvement over the unmanaged Skype, Windows Live, and Yahoo Messenger implementations found in many organizations.

UC can provide flexibility when it's needed, and employees are more in control of their schedules.

QoS and CAC in Unified Communications

A major factor in a successful unified communications (UC) rollout is to ensure that you make efficient use of the available bandwidth on your network. You need to provide a good end user experience if you want people to actually use UC.

Quality of Service (QoS) is a mechanism whereby certain traffic on an IP network can be prioritized over other traffic. This is typically used for traffic such as audio and video, which support real-time communication.

Call Admission Control (CAC) provides a mechanism to define bandwidth available over network links between sites. It's then possible to limit the number of video and audio calls to ensure that saturation point isn't reached. Microsoft Lync provides extended capabilities to allow video elements of a multi-modality call to be routed over different networks (WAN or Internet connections) to further optimize network usage. ■

InstantDoc ID 144383

Planning Ahead

Now that you're familiar with the technical aspects of a UC system, you need to understand how to engage various technical teams and the business as a whole in adopting a UC solution. Marketing, training and reporting are especially important in implementing a successful UC project.

One of the principal guidelines is to bring people together early. Instead of falling into the trap of having teams work against each other, introduce people to your UC vision and get them working together. Achieving this goal requires some forethought. You need to understand how the project will affect each group, and you should outline the opportunities it brings them. For example, in organizations in which a dedicated team runs the telephone system, the staff might feel threatened by a move to an IT-based system. Paint the move to UC as an opportunity: Voice skills aren't going away, they're being augmented with additional skills that are more relevant to today's businesses.

Because UC projects have such a significant business focus, it would be ridiculous not to include the business in the project. It's crucial that you engage with all areas of an organization, possibly by creating project boards that consist of representatives from each area. You need to ensure that you don't design the project based on what IT *thinks* the business needs are—instead, you must seek feedback from a variety of sources within the organization. Pay particular attention to the opinions of administrative assistants, receptionists, and HR staff. These people are the lifeblood of the organization. They interact with a wide variety of people and often have senior connections. If they feel that a project is moving in the right direction, they can help foster good feelings about the project throughout the business. In addition, they can provide valuable feedback to help guide the project.

Part of the role of your chosen champions within each department is to help with internal marketing. Marketing an internal technology

rollout to the business as a whole is extremely important. UC has the power to change the way people work—so the staff must be on board with that change. Creating excitement about upcoming improvements is definitely a marketing task.

User Acceptance Testing (UAT), or piloting, is another important part of successfully implementing UC. The IT department shouldn't undertake this task. Instead, you should use the relationships that you created through the steering groups and project boards to get users involved. Hold demonstration sessions that users can attend to get their hands dirty with the technology. You could set up an informal demo in the lunch room, where users could come by to test the technology and provide feedback. Or you could roll out a more formal testing program, in which you deploy the technology to a particular department and solicit user comments. The key is to use as broad a range of participants as possible and to carefully monitor and address feedback. You need to clearly outline the goals of the test and communicate how the project relates to the overall business strategy. The bigger buzz you create about the project, the more helpful your users will be in testing it.

After the testing phase comes the training phase. Because training can take many forms, it's important to ensure that the correct people get the appropriate amount of support and have access to the necessary resources. Some users will simply need to make and receive phone calls, whereas some will need to manage calls for others, create and lead conference calls, and generally work with a broad range of the modalities provided. The first group might need only some self-training, through walkthroughs or quick tips accessible on the corporate intranet. Power users, however, might need specialty training. Of course, if you've run the project well from the beginning, you already identified those key people early in the process, and they've been involved in training since the testing phase.

In addition to task-based training, you need to train users on etiquette. Over the past 10 or 15 years, people have become accustomed

**True UC integration
requires a cultural
shift within the
organization.**

to how email works and have learned the customs associated with it, such as replying to all, using a signature, and creating out-of-office messages. UC in general and Lync in particular provide entirely new ways to communicate. One concept that might be new to many users is that of presence (i.e., availability). Employees need to know how to use presence within their organization. For example, someone might be in a meeting but available to answer a quick IM. You need to let users know that if they're truly unavailable, they should change their presence to Do Not Disturb. In addition, it's important for users to assign their contacts to the appropriate access group level, so that only those who they want to be able to break through the Do Not Disturb state can do so. Finally, users should be trained to start an IM with "Do you have a moment to talk?" rather than immediately barging in with a direct question.

Following training, it's useful to have some type of accreditation. Many companies give users credit for completing training, and UC training should be no different. You could have users complete an online quiz on a company portal, or you could administer a specific test after training. The method of accreditation you choose will depend on the culture of your company, but it's worthwhile to measure users' capabilities after training and to reward new education.

Regardless of how much marketing, communication, testing, and training you do, some people are simply averse to change and won't embrace it. Proper monitoring can help track general usage to determine areas of the business that aren't fully utilizing the new technology. This knowledge can help you direct another round of communication and training if necessary. In addition, monitoring is helpful to ensure that your solution is performing adequately and to help prove ROI.

If you convince users of the new functionality's relevance to their jobs, you'll have higher acceptance and utilization, which can in turn positively affect the culture of your business. UC training might cost more per employee than a typical IT project's training, but the rewards can also be that much greater.

Culture

For a UC rollout to truly be successful, adoption must occur naturally. You need to create a buzz through marketing and communication that spreads virally throughout the business. The key is to highlight the technology's relevance to all roles within the organization. A simple way to do this is to explain the flexibility of being able to communicate from anywhere that you have an Internet connection. In some cases, explaining the UC solution's importance to the organization requires more insight into the business's goals and processes. More complex examples might include enabling rapid signoff of project documentation through intelligent routing to available managers or integration of Lync to Customer Relationship Management (CRM) systems for Help desk agents to let them quickly see who is calling, view the outstanding customer history, and seek backup from available experts while still on the customer call.

These scenarios illustrate the UC technology's benefit to the business, including cost savings that result from travel reduction and increased use of conferencing. Of course that doesn't mean that you shouldn't continue to meet with people personally; however, after your initial face-to-face interaction, you might follow up with an audio or video conference.

Another example of a situation in which UC would be extremely helpful is if certain meeting participants play only a minor role in the discussion and can dial in to the meeting rather than attend in person. This scenario will be familiar to technical folks—although the salespeople need to build face-to-face relationships with clients, technical staff can dial in to calls, thereby saving travel costs and allowing them to attend other meetings the same day. Keeping your highly trained technical resources in-house as much as possible helps increase your organization's efficiency.

One of the biggest changes that UC brings about is that it can provide a flexible work environment. Although many businesses are reluctant to trust employees, empowering your staff in this way is

beneficial to worker morale. UC technology allows employees to work from anywhere, anytime. Workers are empowered when they're trusted to carry out their roles no matter where they are. Some people might argue that this "flexibility" actually results in constant work. I prefer to look at the benefits—UC can provide flexibility when it's needed, and employees are more in control of their schedules. Of course, giving users this level of trust requires HR support—and not all roles within an organization are flexible. Performance reviews, regular check-ins with managers, and career development continue to be important.

Understanding the cultural change that UC brings is important in fighting a common misconception that can arise in Lync telephony projects. Although Lync can replace a telephony system or conferencing platform, Lync isn't simply a replacement for whatever was previously in place. Lync can take over many functions previously provided by other communication platforms; however, Lync is much more. It's a unified platform that allows not only communication in its own right but also communication as part of the business process.

It's essential that the highest-level executives in a company have the correct mindset regarding Lync. If your senior personnel think of Lync as "just another phone system" (or IM client), your ability to effectively implement Lync throughout the company will be limited. Use all the resources available from Microsoft and the partner community to demonstrate the wider benefits and uses of Lync. (For information about the resources that are available, see the sidebar "Microsoft Resources for Lync Adoption and Training.")

Embrace Change

Deploying a UC solution is a major undertaking if you do it correctly. You could roll out UC just like any other IT project—however, you'll get far more from the technology if you fully integrate it into your business. True UC integration requires a cultural shift within the organization.

Microsoft Resources for Lync Adoption and Training

To successfully adopt Microsoft Lync as part of your unified communications (UC) solution, you need to market it within your organization and train your staff to use it effectively. Microsoft provides numerous resources to assist with your Lync rollout.

- [Lync Adoption and Training - Awareness Resources](#)—Resources to help you promote Lync and communicate with users as you pilot and roll out your solution; includes templates for posters, T-shirts, and a variety of emails.
- [Lync Adoption and Training Downloads](#)—Various user education and training resources for Microsoft Lync, including What's New articles and videos, Quick Start guides, Work Smart guides, videos, and PowerPoint training presentations.
- [Microsoft Lync 2010 Adoption and Training Kit: Lync Custom Intranet Site](#)—Combines various resources, including HTML Help pages and videos, into a customized intranet site that you can deploy within your organization.
- [Microsoft Lync 2010 Adoption and Training Kit](#)—Complete kit, including sample client application add-ins.

These materials can be customized to fit your business needs. For example, you can modify the web pages and publish them on your company's intranet, make the pages available for local download, distribute information through individual Word documents, or send periodic quick-tip email messages. You can repurpose any of the content for either structured training or self-help. ■

InstantDoc ID 143992

You need to empower users to take advantage of the flexibility that UC offers. The technology gives you the tools; all you (and your employees) need to do is embrace it. And if you change the culture of your business, you just might change your business as well. ■

InstantDoc ID 143991

10 Steps to Migration: Configuration Manager 2012

New migration tools simplify the move from SCCM 2007



Peter Daalmans

is a senior technical consultant for a Microsoft Gold Partner, IT-Concern B.V., based in the Netherlands. His specialties are Exchange Server and System Center. He is a Microsoft MVP in System Center Configuration Manager.

Email



Twitter



LinkedIn



Blog



Microsoft's release of the System Center 2012 suite at the Microsoft Management Summit in Las Vegas this April included System Center Configuration Manager (formerly known as SCCM), which takes care of client, server, and device management. Configuration Manager 2012 is all about "the new way of working": Bring Your Own Device (BYOD) and work anywhere. The new version of Configuration Manager takes a big step in that direction by focusing on the user instead of the device. Configuration Manager 2012 is rebuilt on three pillars: empowering users, unifying infrastructure, and simplifying administration.

The first pillar represents the ability for users to access their applications on any device. The second pillar represents the integration of other System Center products, such as System Center Endpoint Protection, parts of System Center Mobile Device Manager, and support for Microsoft Application Virtualization (App-V). The third pillar represents the simplified hierarchy of Configuration Manager 2012 and its support for roles-based administration. You can map your organizational roles to administrative roles in Configuration Manager 2012.

User-Centricity

Microsoft has transformed Configuration Manager from a systems-management platform to one that centers on the user. To support the user-centricity principle, Microsoft developed a new application model. This model uses three new terms, as Table 1 shows.

Table 1: Application Model Terminology

SMS and Configuration Manager 2007 Term	Configuration Manager 2012 Term
Package	Application
Program	Deployment type
Advertisement	Deployment

An application contains data about the application and can hold one or more deployment types. A deployment type specifies the installation files and installation method, such as one of the following:

- Windows Installer
- App-V
- script
- Nokia Symbian
- Windows Mobile installation

Other methods are planned:

- Citrix XenApp application (by Citrix)
- Deep link to a Windows 8 Metro source (SP1)
- Apple iOS applications (SP1)
- Apple Mac OS X applications (SP1)
- Google Android applications (SP1)
- Windows Phone applications (SP1)
- App-V 5.0 applications (SP1)

For each deployment type, you need to set requirement rules and settings. These rules will be used to determine how applications need to be installed on a device on which the user is logged on. The application can be deployed to a user collection; depending on the requirement rules, a deployment type can be used to install or start the application.

Do we need to recreate all our old packages and programs in Configuration Manager 2012 and its new application model? No, because Configuration Manager 2012 has a migration feature.

The Migration Feature

Until now, there hasn't really been a manageable way to upgrade older versions of Configuration Manager (e.g., upgrading Microsoft Systems Management Server—SMS—to SCCM 2007). The most common method was rebuilding the environment from scratch and then recreating the objects by hand, or by upgrading the environment in place if the hardware was suitable for the new version.

The migration feature is an important new feature in Configuration Manager 2012. This feature allows you to preserve the time and money you've invested in SCCM 2007 and migrate it to the new Configuration Manager environment. The migration feature even lets you migrate the environment in a phased way instead of in one big bang.

The process of migrating to Configuration Manager 2012 is a side-by-side migration, so you need to build a new Configuration Manager 2012 hierarchy next to the current Configuration Manager 2007 hierarchy. The migration feature helps you with the migration of most objects through three types of migration jobs:

- collection migration
- object migration
- previously migrated object migration (sometimes referred to as the *objects modified after migration* job)

The collection migration job is used to migrate selected collections with all related objects, such as a package, advertisement, or configuration baselines. You can use the object migration job to migrate selected objects. The previously migrated object migration job is used to migrate objects that were previously migrated but have since been updated in the SCCM 2007 source hierarchy. You can use these migration jobs to migrate these types of objects (as Figure 1 shows):

- advertisements
- App-V packages

- Asset Intelligence catalog
- Asset Intelligence hardware requirements
- Asset Intelligence software list
- boundaries
- collections
- configuration baselines
- configuration items
- OS deployment boot images
- OS deployment driver packages
- OS deployment images
- OS deployment packages
- software metering rules
- software packages
- software update deployment packages
- software update deployments
- software update lists
- task sequences

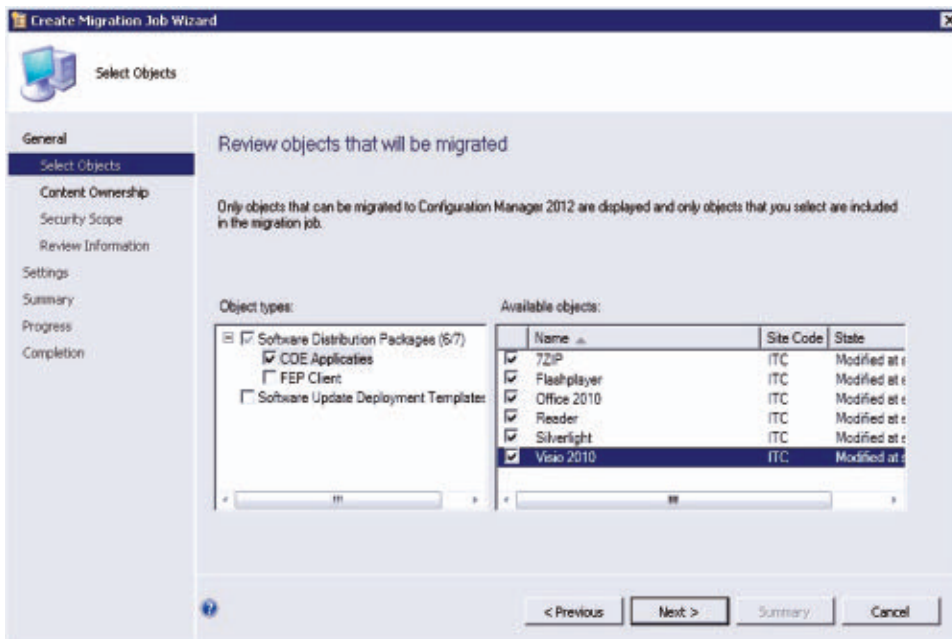
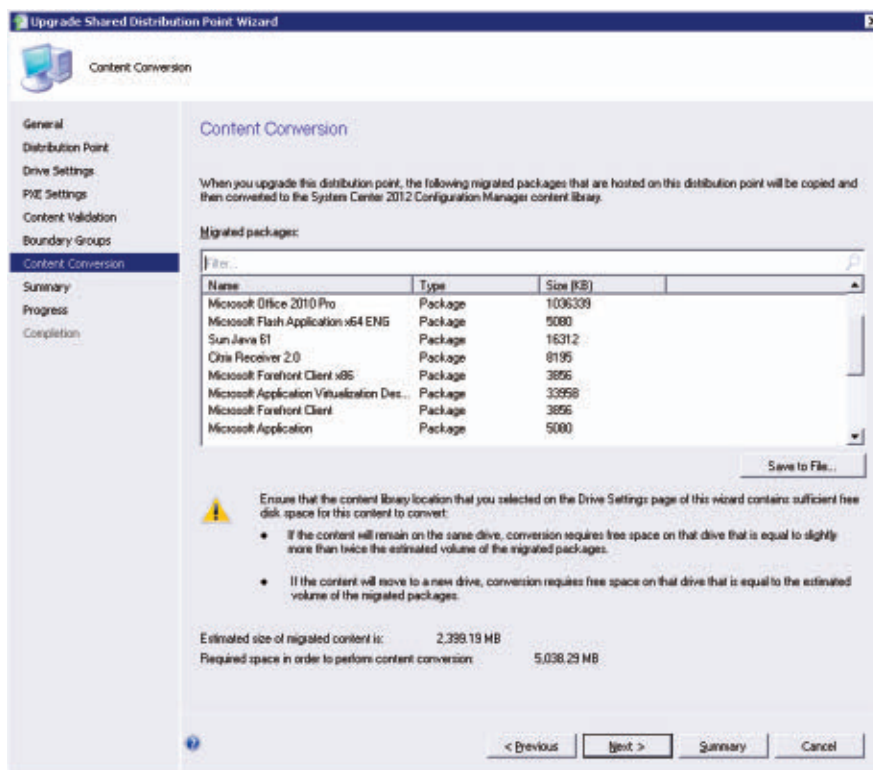


Figure 1
Remigrating Changed
Classic Packages

During the migration process, you can share distribution points that still reside in the Configuration Manager 2007 environment with your new Configuration Manager 2012 clients. This way, your migrated clients can access content on the old distribution points during the migration process. After migrating all your objects and clients to the new environment, you can migrate a distribution point automatically or manually. The server on which the distribution point resides needs to meet certain requirements if you want to migrate a distribution point automatically (as Figure 2 shows):

- Enough disk space must be available to migrate the content to the new content library. The old content needs to be deleted manually after you determine that the migration has finished successfully.
- Only a management point of a secondary site server role may reside on the server. The secondary site server roles are uninstalled

Figure 2
Converting
Distribution Point
Content to the New
Content Library



automatically when you choose to upgrade a distribution point automatically. Roles such as PXE service points, software update points, or state migration points must be manually removed before you upgrade a distribution point.

You can automatically upgrade distribution points such as branch distribution points, distribution point shares, and standard distribution points. When a branch distribution point is installed in Windows XP, you first need to upgrade the OS to Windows 7. Use the branch distribution point management task to migrate the distribution point from XP to Windows 7 without redistributing the content. A big advantage of the ability to migrate the distribution points is that you don't need to use the WAN because the content stays on the server.

Other Migration Tools

Microsoft has developed two other Configuration Manager 2012 migration-related products: the [Package Conversion Manager](#) and the [Physical-to-Virtual \(P2V\) Migration Toolkit](#). Another helpful tool for easing the migration is the [Coretech Package Source Changer](#), developed by the Configuration Manager community.

Package Conversion Manager. The Package Conversion Manager is a Configuration Manager console add-on that lets you convert your migrated classic packages to the Configuration Manager 2012 application model. Classic packages that are migrated from Configuration Manager 2007 are supported in Configuration Manager 2012 but don't support use of the new model's enhanced features. Using the Package Conversion Manager to migrate classic packages is highly recommended.

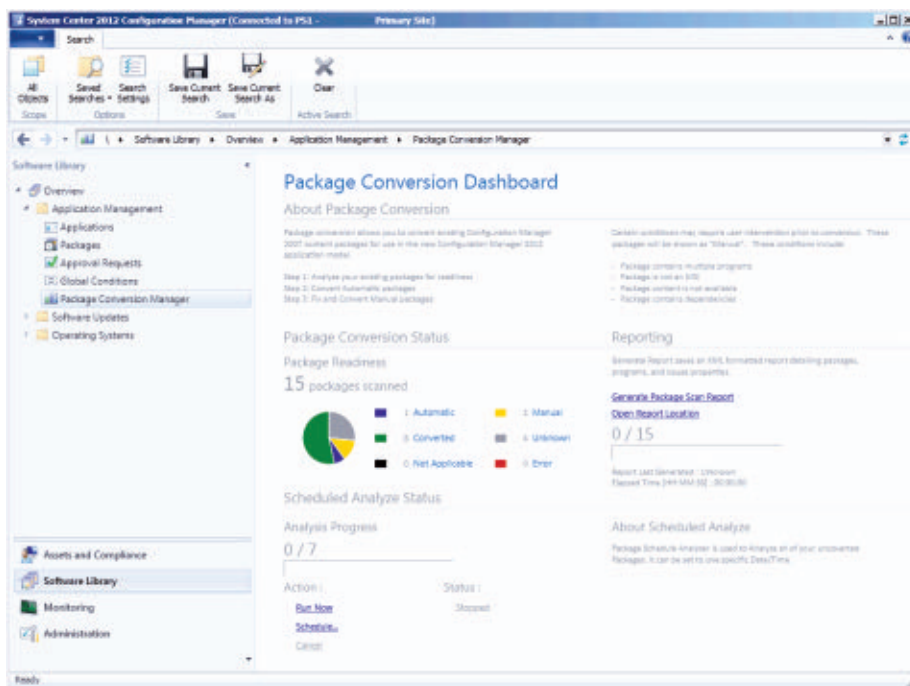
After you install the Package Conversion Manager, which Figure 3 shows, the packages ribbon in the Configuration Manager 2012 console is extended with the following three options:

- Analyze Package
- Convert Package
- Fix and Convert

The Analyze Package option analyzes the classic package and determines whether it can be converted to the new application model. After the classic package has been analyzed, it is assigned a readiness state. Based on its readiness state, a classic package can be converted automatically, manually, or not at all.

To convert a package that has a readiness state supporting automatic conversion, simply select the package and then click the Convert Package button in the Configuration Manager 2012 console. A package with a manual readiness state might have a dependent package that must be converted, or there might be no source content available for the primary site. Depending on the issue, you can use the Fix and Convert option to identify and fix it, with the help of the Fix and Convert Wizard. Before every conversion, the package is analyzed and the validity of the readiness state is determined, as is the package's validity for conversion to the new application model. After the conversion is completed, you still need to test the deployment of the applications

Figure 3
Package Conversion
Dashboard



in your lab environment. Classic packages such as driver software or programs such as defrag are not applicable for conversion to the new model because they are hardware-related instead of user-related.

P2V Migration Toolkit. With the P2V Migration Toolkit, which Figure 4 shows, you can migrate a Configuration Manager 2007 primary site server on a branch office by temporarily virtualizing the server. This step allows you to migrate the site server without investing in new hardware. The P2V Migration Toolkit lets you create a task sequence on standalone media, such as a DVD or USB stick. Or you can create a task sequence that automatically virtualizes your current Configuration Manager 2007 installation by creating a virtual hard disk (VHD), installing Windows Server 2008 R2, configuring Microsoft Hyper-V, and creating a virtual machine (VM) for the old Configuration Manager site server. After this P2V process is finished, you can install Configuration Manager 2012 on the physical server and migrate the virtual Configuration Manager 2007 primary site server via the migration feature. The [P2V Migration Toolkit](#) can also be used as a standalone application, to virtualize a server that is not related to Configuration Manager.



Figure 4
P2V Migration Toolkit
Wizard

Coretech Package Source Changer. The Configuration Manager community is very active, developing numerous scripts and tools as freeware or in the public domain. One such tool is the Coretech Package Source Changer, which was originally developed to change the source paths of packages in Configuration Manager 2007 but which also works with Configuration Manager 2012.

With this freeware tool, you can change the package source of a package object in Configuration Manager and copy the source to a new Universal Naming Convention (UNC) share. This tool is useful for moving or copying the package source of objects that you migrate via the Configuration Manager migration feature.

10 Steps to Migration

The Configuration Manager 2012 migration process is usually part of a project that consists of diverse Microsoft Operations Framework phases. Be sure to plan your migration project and walk through the Envision, Plan, and Design phases. A well-planned and well-designed Configuration Manager environment, based on the results of the envisioning phase, is a must for a successful migration process.

The Configuration Manager migration process consists of 10 global technical steps:

1. Prepare your migration.
2. Test your migration scenario.
3. Configure the migration feature.
4. Configure distribution point sharing.
5. Create migration jobs and migrate the objects.
6. Change the UNC paths of the packages in Configuration Manager 2012.
7. Convert the packages to applications.
8. Migrate the secondary sites and upgrade the distribution points.
9. Deploy the new Configuration Manager 2012 client.
10. Remove Configuration Manager 2007.

Let's look at each step in more detail.

1. Prepare your migration. When you plan to migrate to Configuration Manager 2012, you need to prepare your Configuration Manager 2007 environment to support the migration. You can perform the following steps to add migration support to your Configuration Manager 2007 environment:

- Configuration Manager 2007 SP2 needs to be installed on all site servers; whether you install R2 or R3 doesn't matter.
- There's no support for users and devices in one collection, so create separate collections for users and devices. Collections that reference a collection of a different resource type aren't supported.
- Be sure that your package source is always a UNC path. Local paths will not work when migrating the objects to a new Configuration Manager 2012 server.
- Use unique site codes for your Configuration Manager 2012 environment.
- Upgrade your XP branch distribution points to Windows 7.

2. Test your migration scenario. When migrating your assets from Configuration Manager 2007 to Configuration Manager 2012, it's important to test your migration scenario first in a lab environment. Familiarize yourself with the migration steps in the lab environment before migrating your production environment.

3. Configure the migration feature. You can find the migration feature in the Administration workspace in the Configuration Manager 2012 console. To configure this feature, you need to first define a source hierarchy. This source hierarchy is usually the highest primary site server in the Configuration Manager 2007 hierarchy. After the source hierarchy is defined, a data-gathering process gathers all information about the source hierarchy. When this process finishes for the first time, you need to configure other sites in the hierarchy with credentials that have access to those sites. By default, the data-gathering process runs every four hours.

4. *Configure distribution point sharing.* Not only are all objects inventoried during the data-gathering process, so are the Configuration Manager 2007 distribution points. The eligibility of those points for sharing with Configuration Manager 2012 is also determined. You can configure distribution point sharing per site. Use this option when you have a phased migration of clients, before moving all your packages to a Configuration Manager 2012 distribution point.

5. *Create migration jobs.* Depending on your environment, you can migrate your objects in one or more migration jobs. If your migration phase will take longer, you can remigrate the changed objects in Configuration Manager. Be aware that the longer you are in the migration phase, the longer you need to maintain two Configuration Manager environments.

6. *Change the UNC paths of the packages in Configuration Manager 2012.* When the source of your just-migrated packages is still on the Configuration Manager 2007 site server, you might want to move them over to the new Configuration Manager 2012 server or servers. You can do so by using the Coretech Package Source Changer. This tool can change the UNC path of the package source and copy the files and folder structure to the new package source share on the Configuration Manager 2012 server. By changing the package source at the Configuration Manager 2012 packages, you leave the original source at the Configuration Manager 2007 site intact.

7. *Convert the packages to applications.* After changing the UNC paths and moving the content to the new Configuration Manager site server, you can convert the classic packages to the new application model. The Package Conversion Manager add-on will help you in the conversion process.

8. *Migrate secondary sites and distribution points.* You can't migrate secondary sites to Configuration Manager 2012, so take your time to investigate whether a distribution point can replace your old secondary sites. Most roles, such as PXE support and bandwidth throttling, are now also available when implementing distribution points only.

9. Deploy the new Configuration Manager 2012 client. After preparing and testing the new Configuration Manager 2012 environment, you're ready to deploy the new Configuration Manager clients to your devices. The deployment of these clients can be done in several ways. The best option is to deploy the Configuration Manager 2012 clients with your old Configuration Manager 2007 environment. This way, you have a managed way of deploying the new clients.

10. Remove Configuration Manager 2007. To remove Configuration Manager 2007, you first need to stop the data-gathering process and clean up the migration data from the Configuration Manager 2012 database. After doing so, you can remove Configuration Manager 2007 by uninstalling the site servers.

Migrating to Configuration Manager 2012

This article has given you an overview of the Configuration Manager 2007-to-Configuration Manager 2012 migration approach. You can get started by [downloading the P2V Migration Toolkit and Package Conversion Manager](#) and [Coretech Package Source Changer](#). ■

InstantDoc ID 143642

Microsoft Windows 8 Arrives

The new client OS represents a radical departure from previous Windows versions

Windows 8, Microsoft's latest client OS, features a new UI designed to be tablet touch-friendly, and became available to customers via software upgrades or with new PC purchases on October 26, 2012. Windows 8 represents a radical departure from previous Windows versions and is arguably the most dramatic upgrade Microsoft has yet developed.

The system is essentially a brand-new mobile platform that has been melded onto the traditional Windows desktop, giving users what Microsoft calls a “no compromises” experience that blends the best of mobile with the best of Windows. *Windows IT Pro* brings you ongoing coverage of Windows 8, with in-depth treatment of significant features, breaking news, and analysis. Visit our [Windows 8 page](#) for the latest news and technical features. ■

InstantDoc ID 144099

Windows 8 In-Depth

- ▶ [Welcome to Windows 8](#)
- ▶ [Upgrade from Windows 8 Enterprise Eval? Nope](#)
- ▶ [Windows 8 Review, Part 1: The Desktop](#)
- ▶ [Windows 8 Review, Part 2: You Got Your Metro in My Windows](#)
- ▶ [Windows 8 Upgrade Offer for PC Buyers Goes Live](#)
- ▶ [Start: The Windows 8 Era Begins](#)
- ▶ [Enterprises: Now's the Time to Get Your Windows 8 On!](#)
- ▶ [Installing Windows 8 Enterprise Edition Product Key](#)
- ▶ [Will IT Departments Rush to \(or Away from\) Windows 8?](#)
- ▶ [Q: Is there a version of the Microsoft Assessment and Planning Toolkit that works with Windows Server 2012 and Windows 8?](#)
- ▶ [Q: Why, when I enable .NET Framework 3.5 on Windows 8 and Windows Server 2012, does it connect to the Internet and pull down files?](#)

Windows 8 Features

- ▶ Windows 8 Feature Focus: Start Screen
- ▶ Windows 8 Feature Focus: Multi-Monitor
- ▶ Windows 8 Feature Focus: Tiles
- ▶ Windows 8 Feature Focus: Contracts
- ▶ Windows 8 Feature Focus: Lock Screen
- ▶ Windows 8 Feature Focus: Charms
- ▶ Windows 8 Feature Focus: Snap
- ▶ Windows 8 Feature Focus: Switcher
- ▶ Windows 8 Feature Focus: Back Tip
- ▶ Windows 8 Feature Focus: Start Tip

Windows 8 Tips

- ▶ Windows 8 Tip: Upgrade from Windows 7
- ▶ Windows 8 Tip: Upgrade from Windows XP
- ▶ Windows 8 Tip: Upgrade from Windows Vista
- ▶ Windows 8 Tip: Upgrade from the Release Preview
- ▶ Windows 8 Tip: Customize the Desktop
- ▶ Windows 8 Tip: Customize Live Tiles
- ▶ Windows 8 Tip: Overcoming Library Limitations
- ▶ Windows 8 Tip: New Mice and Keyboards
- ▶ Windows 8 Tip: Protect Portable Storage with BitLocker To Go
- ▶ Windows 8 Tip: Customize the Start Screen
- ▶ Windows 8 Tip: Master Keyboard Shortcuts
- ▶ Windows 8 Tip: Manage Notifications

www.windowsitpro.com/windows-8

Product News for IT Pros



ADAssist Turns Your iPhone and iPad Into an AD Management Tool

BitsAbout announced Active Directory Assist 2.5 for iPhone, iPad, and iPod touch. With ADAssist, Windows administrators can manage Active Directory (AD) directly from their mobile devices. ADAssist is ideal for searching and viewing users, computers, groups, and contacts; managing user tasks such as password resets; performing essential account management tasks such as account unlocking; managing the addition and removal of members from groups; and generating quick reports on locked accounts, soon-to-expire passwords, and soon-to-expire accounts. “ADAssist presents an easy and intuitive way to securely manage AD,” said Sami West, IT director at BitsAbout. “It’s not just productive but also intuitive and beautiful.” ADAssist provides end-to-end, secure, and encrypted access to AD over Wi-Fi or cellular networks. It uses the industry-standard Kerberos 5 protocol. The solution is completely agentless, with no software installation or changes required within the enterprise network. [Download the free ADAssist app from iTunes.](#)



COSprint 3.0 Provides Output Management and Print Spooling

Skybot Software announced COSprint 3.0, output-management and print-spooling software for Windows, UNIX, and Linux servers. This release delivers performance improvements with a new ISAM database and simplified architecture. Customers sending many print requests per second will see the most benefit from COSprint 3.0. “In some cases, these changes have resulted in a 100 percent increase in performance,” explained Pat Cameron, director of automation technology at

Skybot Software. “Even platforms with slower CPUs and less aggressive product usage will see at least a 10 percent increase.”

The most notable improvements of COSprint 2.9 (released in January 2012) were support for failover printers, the ability to print to multiple destinations, and improvements to the audit logs. For more information about COSprint 3.0, visit the [Skybot Software website](#).

NovaStor Offers Server Backup Software for Windows Server 2012



NovaStor announced immediate support for Windows Server 2012 by NovaBACKUP software products. NovaBACKUP is a complete family of products that offers comprehensive and easily scalable data protection for businesses of every size and need. From a single file server to a sophisticated network of application servers, virtual machines (VMs), and workstations, NovaBACKUP lets users protect all their critical business data with its line of Server, Business Essentials, and Business Essentials Network products. NovaBACKUP Business Essentials offers comprehensive protection for all the different types of business data distributed throughout the IT infrastructure. NovaBACKUP Business Essentials Network offers centralized management of data protection for your entire network. For more information about NovaBACKUP, check out the [NovaStore website](#).

Embotics Unveils Cloud Automation Functionality



Embotics announced the newest version of Embotics V-Commander for delivering IT as a Service (ITaaS) using both private and public clouds. Embotics V-Commander focuses on transforming the management paradigm from individual virtual machines (VMs) to a wider range of IT services at the business level. End users can now request the provisioning of IT services consisting of logical aggregates of virtual and physical assets from within a single self-service catalog, streamlining the cloud management process to provide IT organizations with increased control, speed, and agility. Embotics

V-Commander offers rapid provisioning, self-service, service catalogs, IT costing and chargeback, workflow automation, resource optimization, and lifecycle management capabilities. Automated request and approval workflows for both virtual and non-virtual IT assets such as cell phones, laptop computers, and physical servers save valuable time and provide more control to IT administrators. For more information about Embotics V-Commander, see the [Embotics website](#).



FalconStor Enables Automated Service-Oriented Disaster Recovery

FalconStor Software announced the latest version of its RecoverTrac technology, offering customers fully automated recovery of complete IT services in any data center environment. Key enhancements in RecoverTrac 2.5 enable lightning-fast recovery in mixed physical and virtual environments, failover and failback between dissimilar hardware, and non-disruptive disaster recovery testing. As a standard feature of FalconStor Continuous Data Protector (CDP) and FalconStor Network Storage Server (NSS) solutions, the RecoverTrac disaster recovery automation tool lets customers protect critical business services and recover these services in minutes. "At the core of every business lies its most valuable asset, data, which must be protected against human error, equipment failure, and natural disasters, while at the same time always being available," said Darrell Riddle, senior director of product marketing at FalconStor Software. For more information about RecoverTrac, visit the [FalconStor website](#).



ExaGrid Adds SecureErase Feature to Security Capabilities

ExaGrid Systems announced that it has extended its security capabilities by introducing a new SecureErase feature to permanently and safely delete confidential data from disk following the backup process. Fully compliant with Department of Defense (DoD) and National Institute of Standards and Technology (NIST) standards, the new feature ensures that organizations can comply when required to remove

backup data from ExaGrid's disk-based backup appliance. For organizations with high security concerns and strict confidentiality rules, a robust feature that securely erases data is critical. SecureErase makes it possible to execute a secure erasure procedure to remove data that must be deleted, while destroying as little of the surrounding data as possible. For more information about SecureErase, check out the [ExaGrid Systems website](#).

Upload Files to Amazon Glacier with CloudBerry Explorer

CloudBerry Lab released CloudBerry S3 Explorer 3.6, an application that lets users manage files in Amazon Simple Storage Service (Amazon S3) just as they would on their local computers. In the new release, CloudBerry S3 Explorer comes with full support for Amazon Glacier, the recently introduced low-cost storage option. Optimized for data backup and archiving, it becomes a supplement to Amazon S3 storage, costing as little as \$0.01 per gigabyte per month. Amazon has claimed that Glacier is as reliable as its original S3 solution but with longer data retrieval time. CloudBerry S3 Explorer users can now access and manage Amazon Glacier storage, create vaults, move data to vaults, and request to download data back to their computers. Users can create vaults in any of the available Amazon Web Services (AWS) regions. For more information about CloudBerry S3 Explorer, see the [CloudBerry Lab website](#). ■



Hard Disk Manager 12 Professional



**Eric B.
Rux**

is a contributing editor for *Windows IT Pro* and the manager of technical support services at a large university in eastern Washington.

Email



LinkedIn



Modern hard disks and OSs are pretty reliable. Generally speaking, they just work. But sometimes things don't go as planned (or designed), and you need to bring out the big guns. Paragon Software's Hard Disk Manager 12 Professional is a very big gun and deserves a place in your tool box. It can help you repair and manage physical and virtual Windows workstations throughout their life cycle.

Installation

Hard Disk Manager comes in a single executable, which you can install on Windows 2000 Professional (Win2K Pro) and later, including [Windows 8](#). The requirements include Internet Explorer (IE) 5.0 or later, an Intel Pentium or compatible 300MHz or higher processor, 128MB of memory, and a hard disk with at least 250MB of free space. Licensing is handled by a product key and serial number combination. Although not installed by default, the installation executable includes the GPT Loader for hard drives that exceed 2TB and the HotCore Driver, which enables Win2K Pro machines to back up and copy locked volumes.

Installation on my Dell Inspiron laptop took only a few minutes. Afterward, I looked through each menu. It didn't take me long to realize that this product is absolutely chock full of useful features.

As Figure 1 shows, the main screen is laid out with the hard disks logically displayed. A variety of information is provided for each one, including the type of file system, volume size, partition size, used disk space, and free disk space. Right-clicking one of the disks brings up a menu of more than 20 tools and options, which is another indicator of just how many features are available in this product. Just from this menu alone, you can

- back up or restore a partition
- burn the partition to a CD-ROM, DVD, or Blu-ray Disc
- format or delete a partition
- move or resize a partition
- convert or “downgrade” the file system to an earlier NTFS version or FAT
- change the cluster size, serial number, or partition ID
- wipe a partition or clear the “free space” that appears empty in the GUI but might have remnant data from earlier file activity
- test the disk surface and check file system integrity
- view and edit disk sectors

I spent some time testing all the product’s features and was impressed with the simplicity that the wizards brought to what could be considered complex tasks.

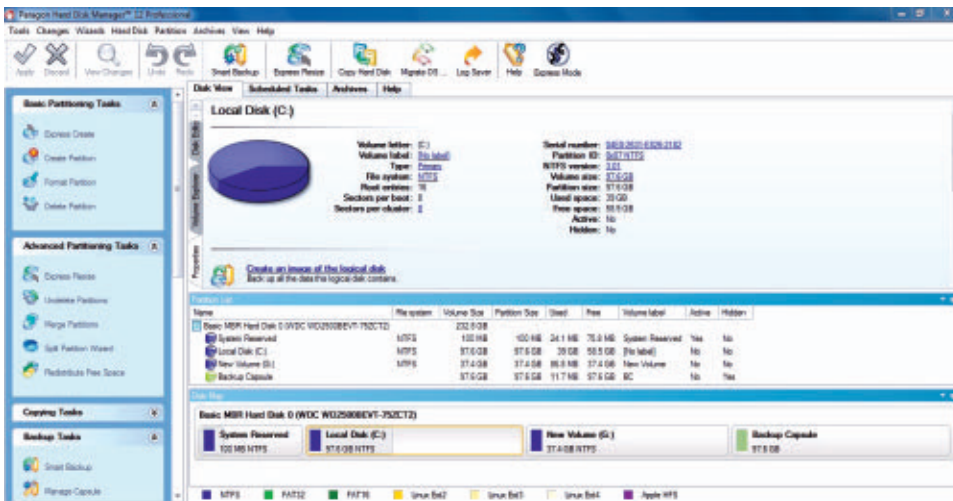


Figure 1
Hard Disk Manager’s
Easy-to-Navigate UI

Data Backup

In an ideal situation, you should have a separate hard disk to back up your important data. This protects you from a total hard disk failure and from “user accidents,” where a file is deleted or changed.

It didn't take me long to realize that this product is absolutely chock full of useful features.

However, most computers are shipped from the factory with only one hard disk. This is where the Backup Capsule comes in. It lets you back up the primary partition to a dedicated place on your hard disk. Although this won't protect you from a total hard disk failure, it will protect you from user accidents. If the primary partition somehow becomes really corrupt, you can even boot from the Backup Capsule by pressing F1 (or another function key of your choosing). Creating a Backup Capsule on a 130GB drive with 40GB of data took about 30 minutes on my laptop and required one reboot.

To protect against total disk failure, you can choose to back up to another local disk, a network server, a CD-ROM, a DVD, or even an FTP server. These backups can be scheduled to occur during system startup, at logon, daily, weekly, or monthly. I found the scheduling to be very feature rich. For example, you can schedule a daily backup with a specific start date, set it to reoccur every X days, and have it stop on a specific date.

Hard Disk Manager also includes classic backup functionality. Called Smart Backups, this feature can protect disks, partitions, email messages (Microsoft Outlook, Outlook Express, or Windows Mail), media files (e.g., photos, videos) stored in users' folders, documents stored in users' folders, and more. You can choose to exclude or include specific file types. The backup wizard can run the backup immediately, schedule it for later, or create a script in the Paragon Script Language. Saving a backup routine in the form of a script makes it easier to transport the backup job from one computer to another.

Partition Management

If you've been working in the IT field for a long time, you'll no doubt remember FDISK. This handy MS DOS utility lets you create and delete partitions on a physical hard drive. One pitfall of this utility was that once the partition was created, it couldn't be resized without first deleting, then re-creating the partition—a procedure that ultimately erased all your data. Hard Disk Manager overcomes this limitation

and lets you move and resize hard drive partitions without losing any data. I successfully carved out a new partition and created a G drive using the built-in wizard.

Other Cool Tools

As I mentioned previously, Hard Disk Manager is full of useful features. Here are some of the other tools that you might find useful:

- **Migrate OS.** Is your hard drive full? Simply purchase a new larger drive, install it into your computer, and use the Migrate OS feature to copy everything to the new drive. Remove the old drive, and you're set.
- **Express Resize.** If one partition has no free space but an adjacent partition has space to spare, you can use the Express Resize feature to move free space from one partition to another.
- **Disk Editor.** For those hard drives that are really messed up, you might need to look at the hexadecimal data or the boot sector. The Disk Editor is an advanced tool that gives you the ability to dig into these areas if needed.
- **Virtual Disk.** I love virtualization for the same reasons everyone else does. But my favorite reason is that it makes disaster recovery so much easier. Virtualization makes disaster recovery a file-restore operation instead of a server-restore operation. The Virtual Disk feature gives you the ability to tweak the virtual hard drives as part of the recovery operation. Microsoft Virtual PC, Parallels, VirtualBox, and VMware virtual drives can all be mounted and treated just like a physical drive.
- **Boot Media Builder.** The Boot Media Builder helps you create Linux or Windows Preinstallation Environment (WinPE) bootable media. This is a separate product installation, but it uses the same licensing product key and serial number combination as the parent product. You must have the Windows Automated Installation Kit (AIK) for Windows 7 installed. A link to it is provided, so finding the right file to download is a snap. Once installed, a wizard

Hard Disk Manager 12 Professional

PROS: Chock full of useful features, including the ability to back up files, folders, partitions, and hard disks; virtual hard drive mounting; raw hex editing

CONS: No undelete capability

RATING: ★★★★★☆

PRICE: \$99.95

RECOMMENDATION: If you work on desktop or server hardware, this is a must have in any PC technician's toolbox. With the exception of an undelete utility, Hard Disk Manager has everything you need to maintain good hard disk health.

CONTACT: Paragon Software • 888-347-5462

walks you through creating bootable media that you can use to start up and recover your PC if the OS fails to boot.

Missing Functionality

With all of the great features in Hard Disk Manager, I was surprised to find that it doesn't include a basic undelete utility. The advent of the Recycle Bin has really helped users have "one more chance" before permanently deleting a file. However, after the Recycle Bin has been emptied or you remove files by pressing Shift + Delete, the files are more difficult to recover. Although there are free applications available that can restore deleted files on your hard disk, why not include this useful feature in a product as robust as Hard Disk Manager? This is a big miss in my opinion and keeps this product from getting a perfect score of five stars. However, even with this feature missing, I would highly recommend Hard Disk Manager to technicians looking for a tool to manage and repair hard disks. ■

InstantDoc ID 143810

Ericom AccessNow 2.0

Before installing a trial version of Ericom Software's Ericom AccessNow 2.0, I checked out the online demos on the company's website. I'm already using Windows 8 Release Preview, so the demos needed to work in Internet Explorer (IE) 10.0. There were no problems on that front. Surprisingly, as I clicked to access a full Windows desktop on Ericom's remote server, I got instant access without needing to install a client-side component. I didn't need an ActiveX control, a Java applet, or any type of client software whatsoever.

Using HTML5, JavaScript, and Cascading Style Sheets (CSS), AccessNow displays remote desktops in web browsers in the same way as a traditional remote desktop client. Although AccessNow doesn't work as quickly as Microsoft's ActiveX-based remote desktop software, the advantages of being able to access a remote desktop or application in any browser that supports web standards, with no client software to install or update, outweighs a minor loss in performance.

AccessNow can connect directly to Microsoft Remote Desktop Services, in which case it relies on whatever authentication mechanism is enabled on the Windows server. Microsoft's Remote Desktop Gateway isn't supported because AccessNow uses HTML5 with WebSockets. However, Ericom includes its own free Secure Gateway server application, which can provide SSL encryption if required.

AccessNow has versions for remote desktop technologies other than those provided by Microsoft, such as Citrix Systems, VMware, and Quest Software. Ericom also has its own connection broker software, PowerTerm WebConnect. It provides options such as centralized management capabilities, reports on who is connected to which servers, and robust enterprise-grade load balancing capabilities. PowerTerm WebConnect also has full support for VMware virtual desktop infrastructures and virtual machine (VM) pools that automatically grow and shrink based on desktop utilization.



Russell Smith

is an independent IT consultant specializing in systems management and security, and author of *Least Privilege Security for Windows 7, Vista and XP* (Packt).



Email



Twitter

Ericom has implemented some clever workarounds to make sure users don't miss out on features found in standard RDP clients.

Installation

AccessNow consists of three parts: the AccessNow server (a WebSocket server), an optional Secure Gateway server, and web components that can be customized to suit your organization's branding and embedded in third-party web pages. Although not required, Ericom recommends that the AccessNow server software be installed on an RDP host for best performance and to enable features such as file transfer. The AccessNow server software is compatible with Windows XP and later.

The AccessNow server software installs in five minutes on your RDP server. The only additional step you need to take is to configure the server's firewall to make sure inbound access on port 8080 is open to AccessNowServer32.exe. Because AccessNow doesn't support Network Level Authentication, I had to change the Windows RDP host connection settings on the Remote tab in the Control Panel System applet to allow connections from computers running any version of Remote Desktop Services (less secure) before I could make a connection through AccessNow.

Once installed, you can connect to the AccessNow server from any client running a compatible browser by entering `http://machine name:8080/accessnow/start.html`, where *machinename* is the name of the server on which you installed AccessNow. You then enter the connection information in the screen shown in Figure 1. Because the AccessNow server is installed on my RDP host, I can leave the first two fields on the connection screen blank. The username and password entered here are passed to Windows for authentication. Clicking Connect took me straight to my remote desktop.

The AccessNow server has a small configuration applet that you can access from the Start menu on the RDP host. From there, you can perform all necessary administration tasks such as managing licensing, restarting the AccessNow service, viewing performance statistics, changing the port on which AccessNow listens for connections, and configuring advanced options (e.g., acceleration, logging, security settings).

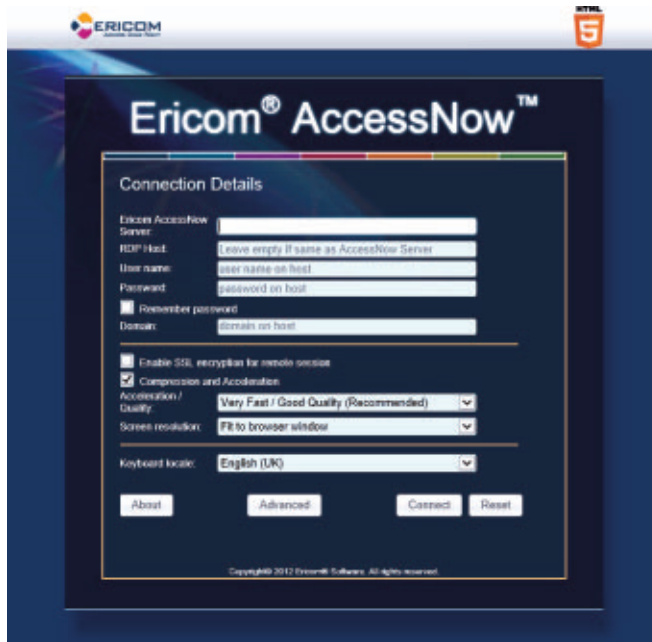


Figure 1
Using AccessNow
to Access a Remote
Desktop in a Browser
Window

File Transfers and Printing

Most browsers impose restrictions on the local resources (e.g., file system, printers) that web pages can access, so Ericom has implemented some clever workarounds to make sure users don't miss out on features found in standard RDP clients. In the top right corner of the browser window, you'll find two buttons for uploading and downloading files between the local and remote desktops. The same functionality can also be achieved by dragging and dropping files in the browser window.

When uploading a file to the remote desktop, you're presented with a standard Windows dialog box to choose the appropriate file. After a short delay, another dialog box appears on the remote desktop asking where you want to save the file. When uploading files to a remote desktop, AccessNow displays a progress bar. When downloading files from a remote desktop, the browser takes responsibility for showing the download progress.

Ericom AccessNow 2.0

PROS: Perfect for giving users access to remote desktops from systems where client software can't be installed (e.g., Google Chromebooks, Windows ARM-based tablets)

CONS: Doesn't work with Remote Desktop Gateway, but Ericom includes its own free Secure Gateway application

RATING: ★★★★★☆

PRICE: Concurrent license starts at \$124; per-seat license starts at \$69; monthly subscription also available

RECOMMENDATION: Using AccessNow is an ideal way to provide remote desktop access with HTML5, which is supported by most modern web browsers. If you don't want to manage remote desktop client installations and updates, AccessNow could be a viable option.

CONTACT: [Ericom Software](#) • 888-769-7876 or 201-767-2210

AccessNow installs its own printer on the remote desktop. If you choose to print to this device, a dialog box appears in the browser window on the local computer offering the option to view the document. This opens the printed document from the remote desktop in Adobe Reader in a separate browser window, where you have the option to print it to a locally installed printer.

Security

AccessNow allows secure WebSocket communications by means of a self-signed certificate. Not all browsers support SSL-encrypted WebSocket connections, so users can opt to disable SSL security in the AccessNow client. Ericom provides the free Secure Gateway server application that can be used in conjunction with AccessNow to provide an additional level of security. The Secure Gateway server authenticates users before they access any internal resources and provides always-on secure WebSocket connections between AccessNow clients and AccessNow servers. Secure Gateway comes with a self-signed certificate installed by default, but it also supports trusted certificates.

Secure Gateway servers can be placed in a demilitarized zone (DMZ) and communicate over the Internet on port 443 using WebSockets. A Secure Gateway server must be a member of the domain in which you want to authenticate users and can be installed on Windows Server 2003 or later and requires Microsoft .NET Framework 4.0.

A Simple, Elegant Solution

AccessNow is a simple, elegant solution for providing access to remote desktops through a web browser, without the need for any client software. With support for most of the standard RDP features provided by Microsoft's own client (including file transfers, printing, and even sound), AccessNow should prove suitable for most standard line of business (LOB) applications. ■

InstantDoc ID 143996

Insights from the Industry

Cluster-Aware Updating and WSUS Presentation

At Microsoft TechEd Australia 2012, I presented a session on Cluster-Aware Updating (CAU) and the new features of Windows Server Update Services (WSUS) 4.0 in [Windows Server 2012](#). You can see the [Cluster-Aware Updating and the New Generation of WSUS](#) presentation on Microsoft's Channel 9.

In case you're unfamiliar with CAU, it's a fantastic technology that works with existing patch management solutions on Server 2012 failover clusters. What happens is that CAU detects that an update is required, obtains the update, puts the first node into maintenance mode, transfers that node's workload to another node, applies the update, restarts the node if necessary, and brings the transferred workload back to the original node. CAU then updates the other nodes in the cluster. The process is quite seamless, as you'll see in the demo toward the end of the presentation.

You can configure CAU to occur automatically, or you can trigger a manual update of all nodes in the cluster. This update method takes a lot of the pain away from cluster updates compared with traditional methods of updating clusters.

Besides demonstrating CAU in the presentation, I discuss how CAU works with WSUS and how to deploy WSUS to a computer running Server Core. I also demonstrate some of the new WSUS PowerShell functionality.

—Orin Thomas
InstantDoc ID 144298



Orin Thomas

is a contributing editor for *Windows IT Pro* and a Windows Security MVP. He has authored or coauthored more than a dozen books for Microsoft Press.



Email



Blog



Video

Cluster-Aware
Updating and the New
Generation of WSUS

IT Professionals Report on Moving Applications to the Cloud



B. K. Winstead

is a senior associate editor for *Windows IT Pro*, *SQL Server Pro*, and *SharePoint Pro*, specializing in messaging, mobility, and unified communications.

Email



Twitter



Blog



Yet another report has come out telling us how many companies are moving applications or data to the cloud. In “[Data Security in the Cloud: Who’s Responsible & How Does It Happen?](#)” I wrote about a survey from [Thales e-Security](#) that revealed the percentage of companies that are using or planning to use the cloud for storing sensitive data. Now [Qumu](#), a video platform provider for businesses, has released findings from its [2012 IT in the Cloud Assessment Project](#). Forty-four percent of the surveyed IT professionals reported that they would be moving applications to the cloud within the next 12 months, and 44.9 percent said they’re already running some applications in the cloud.

One surprising survey result was that more than 30 percent of the IT professionals cited better security as a leading reason for moving to the cloud. Not so long ago, cloud security was the number one reason why IT professionals refused to move applications to the cloud.

According to the survey, the most popular type of application to run in the cloud is email, with 25.9 percent of respondents using the cloud for email services. Email has long been touted as an easy win for moving to the cloud because although any business would rate email as mission critical, running email applications isn’t the focus of most businesses. Other high-scoring application types in the survey were storage (24 percent), document management (13.9 percent), project management (11 percent), and customer relationship management (CRM—10.3 percent).

The survey also reported that 10.3 percent of the surveyed IT professionals outsourced video communication applications to the cloud. With an increasing desire for better collaboration in business combined with smaller travel budgets, it would seem like this category stands to rise dramatically. It’s unclear from the information Qumu released if this category includes conferencing services such as Citrix GoToMeeting or

Microsoft Lync Online, but even if it doesn't, it seems clear that more video is likely becoming a part of all sorts of communications.

Qumu provides a hosting platform for company-made videos as well as a variety of tools for making high-quality videos. The company could be described as “a secure YouTube-like service for enterprise video sharing” but with a suite of products for video creation that can help companies produce branded, useful training and other informative content, and present it to the correct members of the organization.

—B. K. Winstead

InstantDoc ID 144274

Choosing an OS for Exchange Server 2013

One of the interesting decisions awaiting those who want to deploy Microsoft Exchange Server 2013 is the OS to use. You can use either [Windows Server 2012](#) or Windows Server 2008 R2 SP1. Do you choose between Server 2008 R2's well-proven record or Server 2012's promises?

An obvious influence on the debate is the way that Microsoft now treats upgrades for new versions of Exchange. Although build-to-build upgrades are supported within a specific version of Exchange, you haven't been able to upgrade one version of Exchange to a newer release ever since Microsoft discovered the joys of forcing customers to deploy new servers for Exchange Server 2007.

Not surprisingly, Microsoft had a great cover story for Exchange 2007. The complexities of moving a server running a 32-bit version of Windows and a 32-bit version of Exchange to 64-bit versions of the OS and email software were just too horrendous to be contemplated. Panic would ensue if Microsoft even attempted such a feat, so Microsoft basically said, “We're making this really easy for you—buy some nice new 64-bit hardware and have a nice day.”



Tony Redmond

is a senior contributing editor for *Windows IT Pro* and the author of *Microsoft Exchange Server 2010 Inside Out* (Microsoft Press).



Email



Twitter



Blog

There are no technical reasons why Windows Server 2008 R2 can't continue to be the OS of choice for an Exchange Server 2013 deployment—but this statement is also valid for Windows Server 2012.

Ever since then, we've experienced easy upgrades while making it possible for the friendly representatives of our preferred server vendors to make their quarterly sales target. From Microsoft's perspective, the decision to avoid in-place server upgrades makes life much easier for the engineers who maintain Exchange's setup program as well as eradicates a whole heap of potential costs that would otherwise be necessary to code around all the potential "what if" cases involved during in-place upgrades.

One of the obvious difficulties that in-place upgrades would run into is a need to upgrade databases. In the early days of Exchange, this problem didn't exist because Microsoft didn't upgrade the database schema or internal layouts. Roughly the same Jet database persisted from Exchange Server 4.0 in 1996 to Exchange Server 2003 in 2002. Sure, there were tweaks to accommodate structures such as storage groups (SGs), but there wasn't the kind of fundamental overhaul that we've seen recently. Databases are now "upgraded" to a new version by moving mailboxes to servers running the new software. This is a pragmatic and effective method of populating databases. It avoids any need to run an in-place database and schema transformation that might involve hundreds of gigabytes of data. You probably know how long the Eseutil utility takes to rebuild a database. Now contemplate how long it might take to install a new version of Exchange if all the mailbox databases on a server had to be upgraded at the same time.

If you're going to use new server hardware, there's a natural attraction in selecting Server 2012 as the preferred platform for Exchange 2013. After all, it's bright, shiny, and new—and Jeffrey Snover (the lead architect for Windows Server) says that Server 2012 is very good indeed. All in all, given that most Exchange 2013 servers can be expected to be in production until the time comes to deploy the next major release (let's say Exchange Server 2018, as you might ignore any interim releases), it seems to make a lot of sense to use an OS that won't be 10 years old when Exchange 2013 runs out of steam.

However, Windows 2003 SP2 has done a splendid job of supporting Exchange 2007 and it's close to being 10 years old at this point, so maybe an older OS is best. It's been well debugged and is less likely to run into unforeseen issues that a new OS might face. It's also fair to say that some companies have well-developed operational and management processes built around Server 2008 R2, including third-party software. Considerable work might be necessary to upgrade these processes for Server 2012, which might then delay the deployment of Exchange 2013.

Every company is different. No doubt there will be specific circumstances within your company that will tilt the decision one way or the other. What's for sure is that there are no technical reasons why Server 2008 R2 can't continue to be the OS of choice for an Exchange 2013 deployment—but this statement is also valid for Server 2012. So, it's your choice. Heads or tails? ■

—Tony Redmond

InstantDoc ID 144126



En Garde!

**Jason
Bovberg**

Email 

Twitter 

Website 

Well, here's a pretty savvy way to get the word about your products to the younger generation! Intel has created *IT Manager: Duels*, a simulation game for IT pros. Promising to "tap into the



language and humor of IT culture, establishing engagement through game play," *IT Manager: Duels* is the latest installment in the *IT Manager* series of games, whereby players can build and manage their own virtual IT departments, experimenting with Intel tech solutions. For the first time, *Duels* is a multi-player game, and careful strategic planning is required to manage resources and cope with the events that users throw at each other. We think this kind of approach might just work in today's market. After all, as Intel points out, "Intel's main IT decision-maker audience is notoriously averse to tradi-

tional marketing initiatives." Check out *IT Manager: Duels* at [Intel's website](#).




Figure 1: Clearly!



Figure 2: 46 years?!

Send us your funny screenshots, oddball product news, and hilarious end-user stories. If we use your submission, you'll receive a *Windows IT Pro* Rubik's Cube.



Submit

Search our network of sites dedicated to hands-on technical information for IT professionals.

www.windowsitpro.com

Support

Join our discussion forums. Post your questions and get advice from authors, vendors, and other IT professionals.

www.windowsitpro.com/go/forums

News

Check out the current news and information about Microsoft Windows technologies.

www.windowsitpro.com/go/news

EMAIL NEWSLETTERS

Get free news, commentary, and tips delivered automatically to your desktop.

- Cloud & Virtualization UPDATE
- Dev Pro UPDATE
- Exchange & Outlook UPDATE
- Security UPDATE
- SharePoint Pro UPDATE
- SQL Server Pro UPDATE
- Windows IT Pro UPDATE
- WinInfo Daily UPDATE

RELATED PRODUCTS

Windows IT Pro VIP

Get exclusive access to over 40,000 articles and solutions on CD and via the web. Includes FREE access to eBooks and archived eLearning events plus a subscription to either *Windows IT Pro* or *SQL Server Pro*.

www.windowsitpro.com/go/vipsub

SQL Server Pro

Explore the hottest new features of SQL Server, and discover practical tips and tools.

www.sqlmag.com

Dev Pro

Discover up-to-the-minute expert insights, information on development for IT optimization, and solutions-focused articles at DevProConnections.com, where IT pros creatively and proactively drive business value through technology.

www.devproconnections.com

SharePoint Pro

Dive into Microsoft SharePoint content offered in specialized articles, member forums, expert tips, and web seminars mentored by a community of peers and professionals.

www.sharepointpromag.com

Advertiser Directory

1&1 Internet	1
EMC	2
Metalogix	48
Windows IT Pro	70

Vendor Directory

Adobe	14
Amazon	15, 47, 85
Apple	13
BitsAround	98
Citrix Systems	107
CloudBerry Lab	101
Embotics	99
EMC	37
Ericom Software	107
ExaGrid Systems	100
FalconStor Software	100
GigaTrust	31
Google	15, 47
HP	47
HTC	12

Intel	24, 116
LastPass	25
LifeSize	75
Linux	41
Motorola	13
Nokia	12, 85
Nordic Edge	24
NovaStor	99
Paragon Software	102
Polycom	75
Qualcomm	14
Quest Software	107
Qumu	112
Rackspace	47
Radvision	75
Samsung	12
Skybot Software	98
Thales e-Security	112
TITUS	31
Verizon	47
VMware	107

Windows IT Pro